

DOCUMENT RESUME

ED 429 572

IR 019 533

TITLE Computer Network Security: Best Practices for Alberta School Jurisdictions.

INSTITUTION Alberta Dept. of Education, Edmonton.

ISBN ISBN-0-7785-0350-X

PUB DATE 1999-02-00

NOTE 130p.; A publication of the School Technology Task Group.

AVAILABLE FROM Learning Resources Distributing Centre, 12360-142 St., Edmonton, Alberta, Canada T5L 4X9; Tel: 780-427-5775; Fax: 780-422-9750; Web site:
<http://ednet.edc.gov.ab.ca/technology/>

PUB TYPE Reports - Evaluative (142)

EDRS PRICE MF01/PC06 Plus Postage.

DESCRIPTORS Check Lists; *Computer Networks; *Computer Security; Computer Software; *Computer Uses in Education; Educational Administration; Educational Practices; Elementary Secondary Education; Foreign Countries; Glossaries; Information Industry; *Information Policy; *Information Technology; Internet; Models; Policy Formation; Program Implementation; School Districts

IDENTIFIERS Alberta; Client Server Computing Systems; *Computer Industry; Computer Operating Systems; Connectivity; Remote Access; Technology Integration; *Technology Plans; Web Sites

ABSTRACT

This paper provides a snapshot of the computer network security industry and addresses specific issues related to network security in public education. The following topics are covered: (1) security policy, including reasons for establishing a policy, risk assessment, areas to consider, audit tools; (2) workstations, including physical security, protecting workstation components, and computer viruses; (3) the local network, including the OSI (Open Systems Interconnection) reference model, protocols, network segmentation, network management, network sniffing, and data encryption; (4) servers, including UNIX and other server operating systems; (5) remote access, including technologies, remote access servers, protocols, and authentication/authorization; (6) crackers and hackers, including threats and hacking tools/techniques; (7) Internet firewalls, including functions, issues and problems, types, rules, logs, firewall accessories, buying a firewall, and firewall administration; and (8) applications, including e-mail, directory services, the World Wide Web, and single sign on. Each section highlights unique requirements of school jurisdictions and contains a list of relevant web sites. A glossary of terms is provided, and appendices include descriptions of Windows NT Workstation and Server, Novell NetWare, and Apple Macintosh, as well as a security checklist and evaluation form and a list of related Alberta Education resources. (AEF)

* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *

COMPUTER NETWORK SECURITY

Best Practices For Alberta School Jurisdictions

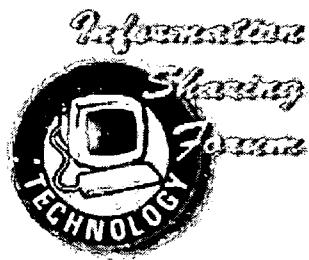
February, 1999

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION
CENTER (ERIC)

This document has been reproduced as received from the person or organization originating it.

Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.



"PERMISSION TO REPRODUCE THIS
MATERIAL HAS BEEN GRANTED BY

C. Andrews

Alberta
EDUCATION

TO THE EDUCATIONAL RESOURCES
INFORMATION CENTER (ERIC)."

ALBERTA EDUCATION CATALOGUING IN PUBLICATION DATA

Alberta. Alberta Education.

Computer network security : best practices for Alberta school jurisdictions.

Available on the Internet: <http://ednet.edc.gov.ab.ca/technology/>

ISBN 0-7785-0350-X

1. Computer networks - Security measures - Alberta. I. Title.

TK5105.59.A333 1998

371.39445

Additional copies are available through:

Learning Resources Distributing Centre
12360-142 Street
Edmonton, Alberta, Canada T5L 4X9
Telephone: 780-427-5775
Facsimile: 780-422-9750

For more information, contact:

Bonnie Brooks
School Technology Task Group
Alberta Education
11160 Jasper Avenue
Edmonton, Alberta, Canada T5K 0L2
Telephone: 780-427-9001
Facsimile: 780-415-1091

To be connected toll free outside Edmonton, dial 310-0000.

The primary intended audience for this framework is:

<i>Administrators</i>	✓
<i>Counsellors</i>	
<i>General Audience</i>	
<i>Information Technologists</i>	✓
<i>Parents</i>	
<i>Students</i>	
<i>Teachers</i>	

Copyright © 1999, the Crown in Right of Alberta, as represented by the Minister of Education.
Alberta Education, School Technology Task Group, 11160 Jasper Avenue, Edmonton, Alberta,
Canada, T5K 0L2.

Permission is given by the copyright owner to reproduce this document, or any part thereof, for
educational purposes and on a nonprofit basis.

PREFACE

Computer network security is a complex technical issue. Careful research, planning, and implementation is required to ensure that an organization has undertaken "reasonable precaution" with respect to the security of confidential and private electronic information.¹

This paper is an introduction to the evolving network security industry. The intent is to provide a current snapshot of this industry to ensure that public educators address security requirements in an ever-changing technical environment. References to Internet web sites and other documents are provided to encourage additional independent research.

Although somewhat technical, the document also is targeted at senior administrators, who also must have a basic understanding of the issues and recommended solutions. The executive summary, sample network schematics, and security checklist are provided to promote this understanding.

¹ The question of "reasonable precaution" also is related to concerns about freedom of information and protection of privacy (FOIPP). FOIPP issues are addressed in another document in this series, *FOIPP and Technology: Best Practices For Alberta School Jurisdictions*.

ACKNOWLEDGEMENTS

PROJECT DIRECTOR

John Percevault,
B. Ed., M. Ed. Director of Technology and Communications
Grande Yellowhead Regional Division No. 35

WITH THE SUPPORT AND ASSISTANCE OF:

Gordon Blinston P. Eng, Ph. D.	Data Communications Consultant TNC The Network Centre Ltd.
Robert Gusnowski, B. Ed., CNE, MCP	Director of Technology Buffalo Trail Regional Division No. 28
Todd Kennedy, B. Sc.	Senior Computer Technician Pembina Hills Regional Division No. 7
Gary Spence, B. Ed., M. Ed.	Manager of Technology Services Wolf Creek School Division No. 72

PROFESSIONAL READERS

Maurice Hollingsworth, B. Ed., M. Ed. Ph. D.	Technology Co-ordinator Palliser Regional Division No. 26
John Stasiuk	Network Operations Team Leader for Computing and Network Services University of Alberta

SCHOOL TECHNOLOGY ADVISORY COUNCIL MEMBERS DURING DEVELOPMENT

John Travers (Chair)	Alberta Education
Peter Balding	Black Gold Regional Division No. 18
Marika Bourque	Calgary School District No. 19
Bonnie Brooks	Alberta Education
George Buck	Universities Co-ordinating Council
David Burch	Alberta Home and School Councils' Association
Edna Dach	Elk Island Public Schools Regional Division No. 14
Jack Dale	Calgary School District No. 19
Peter Darby	Chinook's Edge School Division No. 73
John Darroch	Palliser Regional Division No. 26
Dale Erickson	College of Alberta School Superintendents
Chris Farthing	Alberta School Boards' Association
Toni Hafso	Edmonton Catholic Regional Division No. 40
Gary Heck	Alberta Education
Harry Henshaw	Lethbridge School District No. 51
Robert Hogg	Alberta Education
Maurice Hollingsworth	St. Albert Protestant Separate School District No. 6
Judi Hunter	The Alberta Teachers' Association
Glen Johnson	Palliser Regional Division No. 26
Doug Knight	Calgary School District No. 19
Catherine Kullman	Edmonton School District No. 7
Norma Nocente	Alberta Education
	TELUS Learning Connection
	Universities Co-ordinating Council

John Percevault	Grande Yellowhead Regional Division No. 35
Archie Pick	Alberta Chamber of Resources
Doug Pudwell	Medicine Hat School Division No. 76
Randy Rudd	Association of School Business Officials of Alberta
Ralph Schienbein	Pembina Hills Regional Division No. 7
Jacqueline Skytt	College of Alberta School Superintendents
Ron Sohnle	Elk Island Public Schools Regional Division No. 14
Gary Spence	The Alberta Teachers' Association
Lonnie Springer	Alberta Education
Mary Stephenson	Wolf Creek School Division No. 72
Barbara Stevenson	Calgary Roman Catholic Separate School District No. 1
Paul Stevenson	Alberta Chamber of Commerce
Arwin van Voorthuizen	Calgary School District No. 19
	The Alberta Teachers' Association
	Horizon School Division No. 67
	Alberta College of Art
	Council of Presidents of Public Colleges and
	Technical Institutes of Alberta

RESOURCE PERSONNEL

John Hogarth	ConsultNet
Peter Wright	University of Alberta

TABLE OF CONTENTS

Executive Summary	1
Section I: Security Policy.....	7
The Basics	7
Reasons for Establishing a Policy.....	7
Risk Assessment	7
Areas to Consider	8
Audit Tools	11
Unique Requirements of School Jurisdictions	12
References.....	12
Section II: Workstations	15
The Basics	15
Physical Security	15
Protecting Workstation Components	17
Computer Viruses	18
Unique Requirements of School Jurisdictions	20
References.....	21
Section III: The Local Network	23
The Basics	23
The OSI Reference Model	23
Protocols	24
Network Segmentation.....	27
Network Management.....	28
Network Sniffing	29
Data Encryption.....	31
Unique Requirements of School Jurisdictions	31
References.....	32
Section IV: Servers	33
The Basics	33
UNIX.....	36
Other Server Operating Systems	38
Unique Requirements of School Jurisdictions	38
References.....	39
Section V: Remote Access	41
The Basics	41
Technologies.....	41
Remote Access Servers	43
Protocols	44
Authentication and Authorization.....	44
Unique Requirements of School Jurisdictions	49
References.....	49
Section VI: Crackers and Hackers	51
The Basics	51
Threats	51
Hacking Tools and Techniques.....	52
Unique Requirements of School Jurisdictions	53
References.....	53
Section VII: Firewalls.....	55
The Basics	55
What is An Internet Firewall?	55
Functions Performed by Firewalls	55
Issues and Problems	57
Types	58
Rules	59

Logs	61
Firewall Accessories	62
Buying a Firewall	63
Firewall Administration	64
Unique Requirements of School Jurisdictions	65
References.....	65
Section VIII: Applications	67
The Basics	67
E-mail	67
Directory Services	70
WWW	71
Single Sign On	73
Unique Requirements of School Jurisdictions	73
References.....	74
Section IX: Closing Comment	77
Glossary of Terms.....	79
Appendix A: Security Checklist and Self-Evaluation	89
Appendix B: NT Workstation and Server	91
Introduction	91
Overview	91
Stand-alone NT Workstation.....	96
NT Workstation in the Workgroup	99
NT Server and the Single Server Domain.....	101
NT Workstation in a LAN with Administrative and Instructional Functions	104
NT and the Enterprise	105
Conclusions.....	108
References.....	108
Appendix C: Novell NetWare	111
Introduction	111
Novell Security Model.....	113
NDS Management Model.....	115
Client Management and Security	116
Novell Remote Management Utilities	117
Novell and Internal Security Risks	118
Novell and External Risks	119
Addressing Hardware Reliability	119
Summary.....	120
Utilizing Novell in Public Education	120
References.....	122
Appendix D: Apple Macintosh	123
Introduction	123
At Ease.....	124
Security	125
Network Assistant	126
The Network Assistant as a Teaching Tool	126
FoolProof.....	127
MacJanet.....	127
Summary.....	128
Appendix E: Related Alberta Education Resources	129

LIST OF FIGURES

Figure 1:	Typical Education Network	4
Figure 2:	Typical Network Server Room (per school/site)	5
Figure 3:	Typical Network Central Office Server Room	6
Figure 4:	The NTFS Permissions Dialog	94
Figure 5:	The Audit Policy Dialog	95
Figure 6:	The Object Auditing Dialog	95
Figure 7:	The User Profile Dialog	98
Figure 8:	The Share Permissions Dialog	100
Figure 9:	Editing Policies for the Student Group	103
Figure 10:	The Trust Relationship Dialog	107
Figure 11:	Network Security Threats	114
Figure 12:	A Graphical Representation of the NDS Tree	115

EXECUTIVE SUMMARY

In public education, there is currently an extensive move from discrete, multiple local area networks serving school laboratories, school offices, and jurisdiction offices to an enterprise wide area network approach with connections to the Internet.

Discrete local area networks provide a reasonable level of security, as there are limited opportunities for communications or data exchange within or outside of the organization. However, a corporate-wide network connected beyond the organization must not be constructed without a researched, planned, and regularly updated security implementation policy which acknowledges:

- associated risks,
- available technical solutions,
- capital costs and maintenance contracts, and
- day-to-day overhead in the documentation, monitoring, management, and research of networking and security technologies.

Typically, a school network is comprised of student and staff personal computer access. This network is serviced by on-site file and application servers designed and installed to limit an individual's access to specific files, applications, and server/hardware configuration parameters. Careful planning may address many initial security variables. But, given multiple users' day-to-day interaction, personal security practices, and unrestricted or unsupervised network access, such planning is not without associated risks. The risks include:

- unauthorized access to, and disclosure of confidential and private information, or
- modification of, willful damage to, or destruction of corporate databases and files, or
- the willful interference or corruption of network devices and servers with the intent to interrupt network services.

As jurisdictions continue to connect schools to the jurisdiction office to extend central applications and services to school-based users, additional corporate investments must be secured to ensure that authorized access is maintained. The combined investments in electronic file construction and daily maintenance include a wide variety of installed, working applications:

- centralized finance with remote access,
- electronic human resource systems,
- Freedom of Information and Protection of Privacy (FOIPP) records access database,
- centralized student records registry and student transportation demographics,
- school-based student records systems,
- school library catalog and circulation records,

- counseling and psychological assessment files,
- electronic mail,
- Internet web servers,
- employees' corporate electronic files (teacher markbook, word processor, spreadsheet, etc.), and
- students' personal electronic files (generated work).

The education sector is not unique in its requirement to secure electronic records but it may be potentially more at risk than private sector entities. In education, there are many users accessing different personal computers at different times in different locations within a building. This factor alone defines the requirement to control and monitor physical access to sensitive locations and/or personal computers.

As evolving networks expand school jurisdictions' capabilities to reach into the community, more jurisdictions are likely to consider extending their hours of service to a growing community of stakeholders. Such initiatives will create additional challenges regarding security. Extended stakeholders may include evening and weekend access at different physical locations such as:

- public libraries,
- non-credit community learning societies, and
- credit region educational consortium.

The rapid growth of the Internet with a jurisdiction's intranet investment has further compounded a significant technical challenge. A noticeable result of this rapid growth is the fast-emerging network security industry with the many, ever-changing security solutions designed to secure data investments and data communication transactions.

The main focus of this paper is to take a snapshot of the computer network security industry, and then document specific issues and target solutions. Internet sites are included in the reference lists to facilitate more detailed reviews of topics such as:

- security policies,
- controlled physical access to server rooms and wiring closets,
- network operating system and desktop operating system security,
- enterprise security management (servers, LAN/WAN technologies),
- network design and changing technologies (Ethernet, routing, switching),
- internal and external firewall requirements,
- computer viruses,
- Internet RFC,
- electronic mail (encryption, digital signatures), and
- network dial-in remote access (employee, vendor support, authentication).

Capital acquisition of required security technology with respective maintenance contracts is one component of a jurisdiction's annual expenditures on computer networking technology which is critical and yet often lacking. Given the risks of connectivity, it is imperative that jurisdictions have staff who undertake careful planning to address the issues presented in this document. They must ensure that the organization is indeed making "reasonable efforts" to keep confidential and private information secure.

Doing "next-to-nothing" is at one end of the continuum, while having "state-of-the art" technology is at the other. There can be an appropriate balance somewhere in between. The acquisition of workable, affordable security technology will require a high level of understanding among jurisdiction planners and technical support personnel. These decisions and investments should not be made independently of the required human resources to document, monitor, manage, and research on a day-to-day basis.

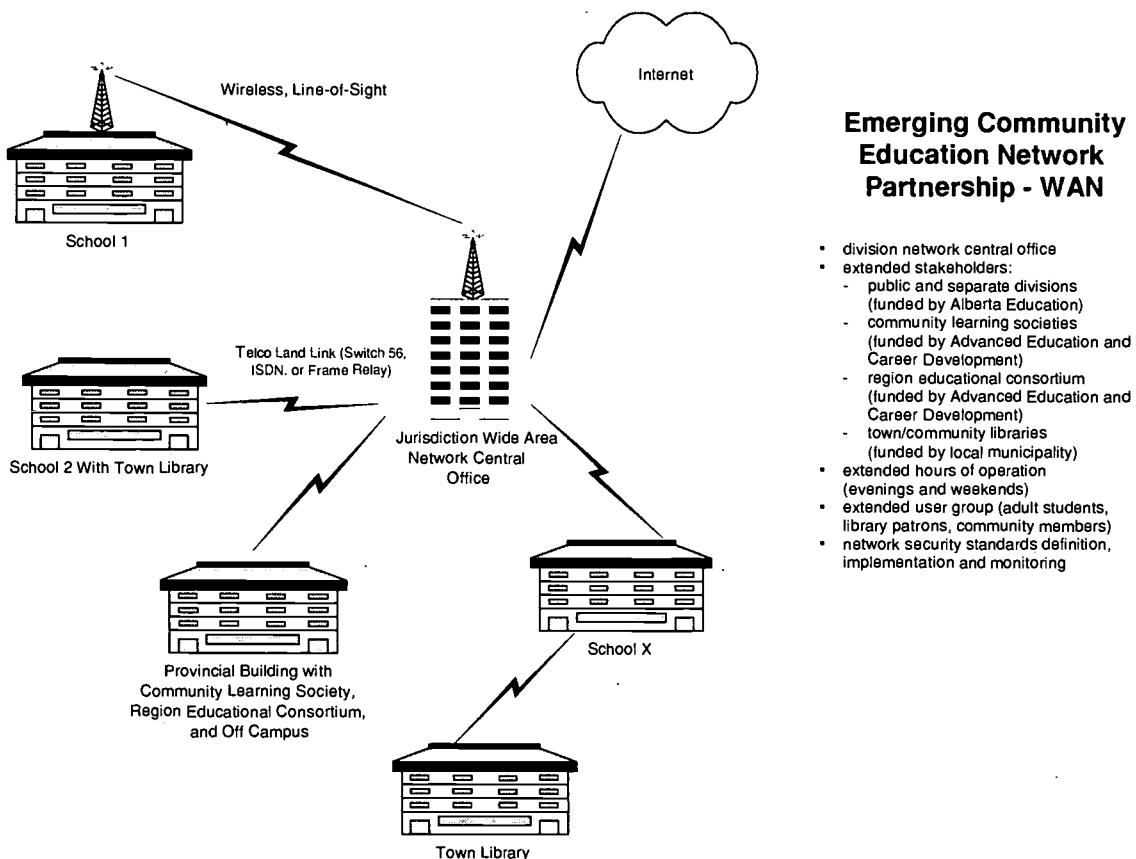
In their quest to become connected and thereby add value to curricular and business processes, public educators must address critical security and total-cost-of-ownership issues. Failure to do so can be interpreted as not taking "reasonable precautions" in securing confidential and private information.

Administrators and managers responsible for network design must consider these key items:

- Technology plans for a school jurisdiction should include a security component with an associated checklist.
- Defining product standards for both hardware and software will limit the number of variables on a network and consequently limit associated security issues.
- For Intel platforms, the preferred, current desktop operating system is Windows NT 4.0 Workstation.
- To gain better bandwidth by collapsing the backbone, switches are recommended over hubs. Switches also provide additional security if one workstation is defined per switch port. This approach limits the ability to sniff the Ethernet segment.
- Routing is the preferred architecture as opposed to bridging. Routing limits many broadcasts to the local area network, not the wide area network.
- School jurisdictions should consider new router designs that support the definition of virtual LANs on a wide area network.
- School system networks need to be protected from the Internet by means of robust, industry-grade firewalls. Sensitive servers on the system wide area network need to be protected from others by means of service networks on the Internet firewall or by the addition of an internal firewall.
- There must be regular audits of network traffic, firewall log files, antivirus log files, and server log files and configurations.
- Limiting physical access to sensitive equipment must be planned from the outset.
- Technology staff development and training must be done regularly to ensure that staff members are up to date on required technical solutions.

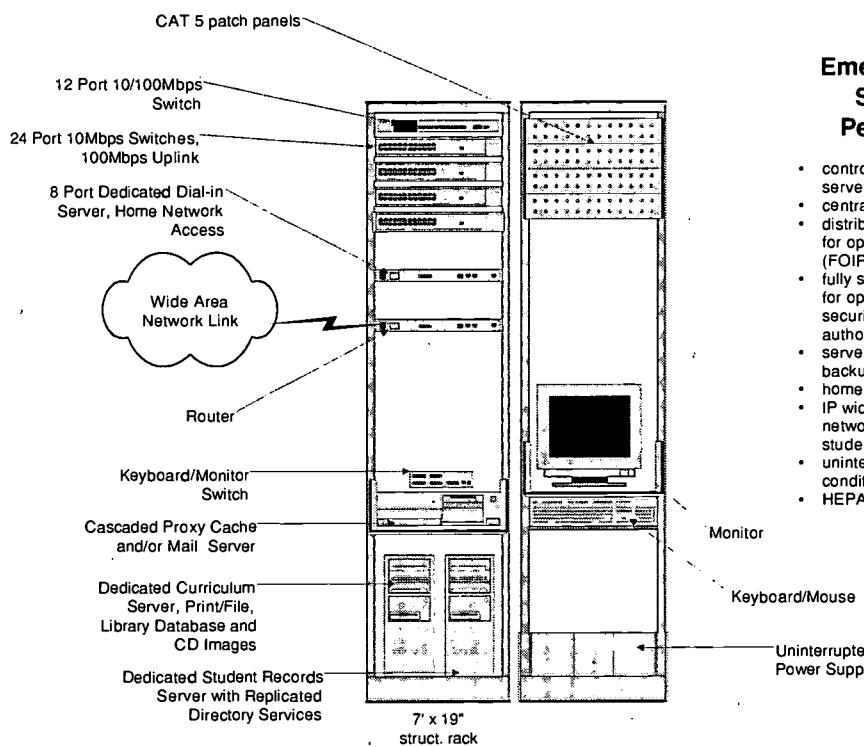
- Inter-jurisdictional communication among technology staff members is an effective way of exchanging timely information in this dynamic industry.

The following diagrams show the characteristics of typical network designs that are emerging in Alberta school jurisdictions.



(Other transport media include satellite, cable, or fibre optics.)

FIGURE 1: TYPICAL EDUCATION NETWORK



Emerging Network Server Room Per School/Site

- controlled physical access—lock down server console
- central termination for CAT5, fibre
- distributed process across multiple servers for optimal network performance and security (FOIPP)
- fully switched, dedicated bandwidth per PC for optimal network performance with packet security, virtual LAN definitions, and authorized port access
- servers with redundant drives and tape backup for data archive
- home dial-in with security authentication
- IP wide area routing (possibly physical network definitions for administrative and student PC access)
- uninterrupted power supply, electrical conditioning
- HEPA air filter

FIGURE 2: TYPICAL NETWORK SERVER ROOM (PER SCHOOL/SITE)

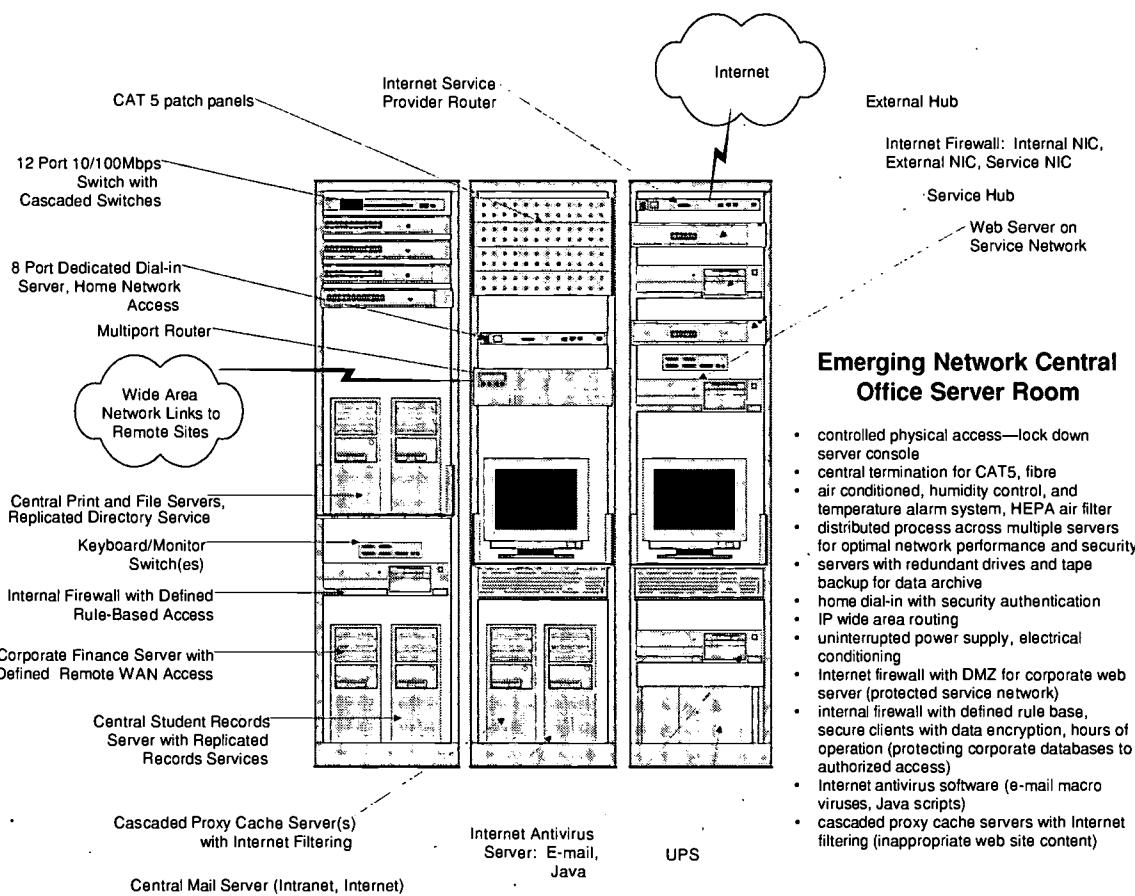


FIGURE 3: TYPICAL NETWORK CENTRAL OFFICE SERVER ROOM

BEST COPY AVAILABLE

SECTION I

SECURITY POLICY

THE BASICS

Various tools are available for detecting potential security weaknesses in networks and servers. Just as it is important to regularly check the door locks in a large facility, it is important to use whatever tools are available for detecting potential security problems.

A security policy is one of the most important components of a comprehensive security strategy. The policy outlines responsibilities for security and provides members of the organization with the information they need to help keep facilities and information secure. The policy is based on an assessment of the security risks, and it includes statements referring to the duties and responsibilities associated with maintaining a secure environment.

REASONS FOR ESTABLISHING A POLICY

An organization establishes policies to let members of the organization know what sorts of activity and behaviours are acceptable and what should be avoided. They allow the members to manage their behaviour and make decisions that are consistent with the goals of the organization.

In most large organizations, the vast majority of computer users may be unaware of the risks associated with data loss or damage to computer systems. They do not know what the potential threats are. Such computer users are like tourists in a strange city who do not know where it is safe to walk or how to deal with strangers who may approach them. If they can find a reliable guidebook that gives them the safety information they need, their vacation will probably be much safer and more enjoyable. A security policy can be viewed as such a guidebook.

The policy writers know about the dangers, and they know the best ways to defend against them. If the policy is well written, if the contents are properly communicated, and if the users follow the rules specified in the policy, they likely will not place themselves or the organization in a vulnerable position.

RISK ASSESSMENT

One of the first steps involved in creating a security policy is the assessment of risk. An assessment involves itemizing all components of the computer system—hardware, software, network, and data—and then deciding what the relative impact would be if any one of these components were damaged or lost. Security measures generally result in greater inconvenience for the users and an increase in cost for the administration. Therefore, the potential effects of loss must be weighed against the effort and cost required to minimize that loss. The cost of maintaining effective security will increase as a jurisdiction extends its reach into the community.

AREAS TO CONSIDER

Since every organization values its assets differently, it is difficult to apply an external set of standards when developing security policy. The guidelines presented below address most major issues, but each jurisdiction will probably have reasons to add to or subtract from this list.

1. *Introduction and Purpose*

- Why are information systems important to the organization?
- What are the technical and business risks due to alteration, theft, inaccuracy, or destruction of data?
- How can each member of the organization be made aware of these risks and motivated to act in a way that protects the organization?

2. *Scope*

- To whom does the policy apply?
- To which computer systems does it apply?
- To which data and software does it apply?
- To which peripheral devices and facilities does it apply?

3. *Definitions*

- This section should contain definitions of all terms used in the policy that might not be clear to the average reader.

4. *Identification and Authentication*

General Identification and Authentication Policies

- Under what circumstances is authentication required?
- What mechanisms are used for authentication?
- Who approves the rights and privileges associated with a particular level of identification?

Password Management Policies

- Who issues user IDs?
- Each user is responsible for all activity which occurs on his or her user ID.
- Under what circumstances will user IDs be revoked?
- What should a user do if s/he forgets the password?
- How often should passwords be changed?
- What criteria should be used when choosing passwords?

Encryption, Digital Signatures and Certificates

- Which applications require the use of encryption, digital signatures, or certificates?
- Which standards should be employed for encryption, digital signatures, and certificates?
- What mechanism should be employed for managing encryption keys?

5. *Software Import Control*

Controlling Software

- What types of software may be installed on workstations and servers?
- What are the approved sources for software?
- Who is responsible for installing software on servers and workstations?

Virus Prevention, Detection, and Removal

- Which workstations will be equipped with antivirus software?
- What antivirus software should be used?
- Who is responsible for insuring that the antivirus software is up to date?
- What is the policy concerning the use of floppy disks that have not been checked for viruses?
- What is the policy concerning virus checking for files or software that has been downloaded from a source outside the organization's network?
- If a virus is detected, to whom should the event be reported?
- What is the process for removing a virus?

6. *System/Architecture*

Remote Access

- Who will be allowed to access the organization's information systems from outside the network?
- What mechanisms will be implemented to facilitate remote access?
- What authentication mechanism will be required for remote users?
- What activities of remote users should be logged?
- What are the restrictions imposed by the Internet firewall?

Access to Internal Databases

- All data files will be protected against unauthorized changes.
- Sensitive data will be protected against unauthorized reading and copying.
- Who should be considered as the owner of data?

- Who specifies the sensitivity level of a file and which user IDs may access or modify it?

Virtual Private Networks

- Under what circumstances will virtual private networks be implemented?
- What controls will be placed on access to virtual private networks?

7. Applications

E-mail

- Which members of the organization have access to e-mail?
- Under what circumstances should e-mail be encrypted?
- What mechanism should be employed to ensure that e-mail is authentic?
- How should e-mail address lists be managed?
- Is e-mail considered to be private, or is it subject to inspection?

WWW

- Is there a policy for restricting access to specific WWW sites?
- If so, who should be governed by the policy?
- Who determines which sites should be restricted?
- Are WWW accesses logged? If so, who has access to the logs?

Student Records

- Who has authorized access to the student records system?
- How is access restricted by user (menu/function, "need-to-know")?
- What mechanism should be employed to limit server access to defined users' systems?
- What mechanism should be employed to restrict students' access to defined users' systems?

Central Finance

- Who has authorized access to the central finance system?
- How is access restricted by user (menu/function)?
- What mechanism should be employed to limit server access to defined users' systems?

8. Incident Handling

Intrusion Detection

- Who is responsible for intrusion detection?
- What monitoring processes should be followed to facilitate intrusion detection?

Incident Response

- Who should be notified if an incident occurs?
- What steps should be taken to recover from an incident?

9. Administration

Assigning Security Responsibility

- Who is responsible for administering the various sections of the security policy?
- Under what circumstances should the policy be revised?

Appropriate Use

- Users should be prohibited from:
 - unauthorized attempts to break into any computer
 - using company time and resources for personal gain, theft, or copying electronic files without permission
 - sending or posting confidential files outside the organization or inside the organization to unauthorized personnel
 - refusing to co-operate with a reasonable security investigation, and
 - sending chain letters through e-mail.

Privacy

- Does the organization expect to be able to inspect any information stored on computer systems or transmitted via the network?
- Are members of the organization aware that their information may be subject to inspection?

Awareness and Education

- Who is responsible for co-ordinating and providing security awareness education?
- What mechanisms will be employed to promote awareness?

Endorsement

- From where must approval be obtained for the contents of the security policy?

AUDIT TOOLS

Security weaknesses are often very obscure, and any evidence that they exist is buried in a flood of legitimate information on a network or within a computer system. To overcome this problem, a variety of computer programs have been written to look specifically for security problems. One of the most highly publicized of these tools is SATAN. When it was released, many people were predicting hackers would misuse it to

find vulnerabilities that they could attack. Undoubtedly, it has been used for this purpose, but it also is a very valuable resource for system managers looking for security weaknesses in their own networks.

SATAN is in the public domain. Similar packages available from commercial sources purport to expand on the capabilities of SATAN and several apply to operating systems other than UNIX.

Another program that hackers use extensively is called CRACK. It attempts to decode UNIX passwords by a trial-and-error method involving a database of potential password choices. A systems administrator can use it to test whether users are choosing appropriate passwords. Similar cracking programs are available for NT passwords.

System administrators can use one other type of tool used by hackers to detect unauthorized modems on the network. This tool automatically dials every telephone number within a particular range and attempts to determine whether a modem answers. One example of such a tool is "toneloc" which is available at
[<ftp://ftp.paranoia.com/pub/tl110.zip>](http://ftp.paranoia.com/pub/tl110.zip).

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

School jurisdictions must ensure that the security policy is communicated and understood by all students and staff. An additional area to consider in the policy is a section that explicitly defines penalties for contravention of the security policy.

REFERENCES

Security Policies

[<http://csrc.nist.gov/policies/welcome.html>](http://csrc.nist.gov/policies/welcome.html)

Audit Tools

[<http://www.net.tamu.edu/policy/tools/intro.html>](http://www.net.tamu.edu/policy/tools/intro.html)

[<http://www.alw.nih.gov/Security/security-prog.html>](http://www.alw.nih.gov/Security/security-prog.html)

[<ftp://ftp.cert.org/pub/tech_tips/security_tools>](http://ftp.cert.org/pub/tech_tips/security_tools)

SATAN

[<http://www.fish.com/satan/>](http://www.fish.com/satan/)

[<ftp://coast.cs.purdue.edu/pub/COAST/tools/>](http://coast.cs.purdue.edu/pub/COAST/tools/)

General Security Information

<<http://www.cert.org/>>

<<http://www.cs.purdue.edu/coast/archive/index.html>>

<<http://www.alw.nih.gov/Security/>>

<<http://www.echotech.com/secursit.htm>>

SECTION II

WORKSTATIONS

THE BASICS

Workstations consist of two separate entities: hardware components such as keyboards, monitors, and cabinets with disks and electronics, and software—the programs that actually control the behaviour of these physical components. Both entities must be maintained in working order and protected from damage.

Protection for hardware involves simple physical security. Protection for software requires physical security as well as some means of preventing damage caused by the direct action of users or more subtle attacks from malicious programs referred to as "viruses." Protection from all of these threats is an essential requirement for any computer installation. Workstations are expensive and easily damaged, and corrupted software can result in a tremendous loss of productivity and a major effort for re-installation.

Various degrees of protection can be achieved with varying costs. The overall expenditure must be balanced against the potential losses. For school jurisdictions, appropriate policies concerning the protection of workstations provide an essential and inexpensive starting point for workstation security. However, such policies are not effective without antivirus software or other appropriate tools.

PHYSICAL SECURITY

Computers are valuable. They are in wide demand. A stolen computer can easily be dismantled and the parts can readily be sold. Also, although modern computers are relatively robust, their design makes them attractive targets for vandals. An unattended computer that has special access privileges to restricted information also is an appealing target for a curious or malicious person. For these reasons, physical security is an important component of a school jurisdiction's overall computer security plan.

LOCATION

The easiest way to protect a computer from physical abuse is to locate it behind securely locked doors. Locked doors are an absolute necessity for all servers and system consoles and should be used wherever possible for any other computer workstations.

Where public access is required, visibility is one of the best forms of protection. Workstations should be located where they can be readily observed, either by staff or some sort of video monitoring system.

TIE-DOWNS

Various devices that can be purchased or built will attach computer workstations to desks, benches, or some other immovable object. These devices have the potential to deter an opportunistic thief, but they also have a tendency to restrict access for maintenance personnel, and they generally do not protect against vandalism. Of course, these devices do not restrict access to the internal components of the workstation.

ALARMS

Electronic alarms that sound a warning if a workstation is tampered with or detached from the network may be useful in particular situations. However, they usually make maintenance more difficult, and false alarms may become a nuisance.

MARKINGS

Indelible markings or non-removable stickers on cases, monitors, and keyboards can deter theft. It is not usually practical to mark each and every internal component of a workstation.

PASSWORDS FOR ACCESS

Various levels of passwords may be established for a workstation. Most workstations allow a password to be entered into the BIOS so that unauthorized persons can not reboot the machine. Screen saver passwords prevent unauthorized access to a running machine that has been left alone long enough for the screen saver to activate. Network operating systems such as Novell NetWare or Windows NT employ passwords to authorize access to individual workstations and network resources, and individual applications such as word processors or e-mail programs allow files to be password-protected.

In selecting and managing workstation passwords, there is a trade-off between security and functionality. Ideally, passwords should be difficult to guess, and they should be changed often. But passwords that are difficult to guess are often not easy to remember; and changing passwords regularly leads to more forgotten passwords and more work.

BIOS passwords are of particular concern. Maintenance personnel must have access to the BIOS on every machine in the organization, but they can not be expected to remember a separate password for every machine. On the other hand, if the same password is used universally, and it is compromised, then every machine in the organization is vulnerable. A horrendous effort will be required to change all the passwords. Because software tools for obtaining and cracking BIOS passwords are available, and because students have direct access to individual workstations, it is advisable to at least use different passwords for student and administrative workstations. As well, if resources permit, passwords should be changed regularly.

Screen saver passwords can be used to prevent access to individual workstations. If possible, all administrative workstations should employ screen saver passwords, and screen savers should be configured to start after a relatively short period of inactivity.

Schools and jurisdictions should have a policy on the format and change intervals for screen saver passwords. Typical policies specify that passwords should be at least six characters in length, and they should contain a mixture of upper- and lower-case alphabetic characters as well as symbols and numbers. Such policies often also specify that previously used passwords should not be re-used.

Network operating systems allow the system administrators to define specific password management policies. They can enforce regular password changing and specify password formats. The management of administrator passwords for specific workstations involves many of the issues associated with BIOS passwords. Maintenance personnel must be able to obtain administrative access to the workstation, but it may be desirable to disallow such access for anyone else.

USER PROFILES

Windows NT allows the establishment of a specific profile for each user who has the capability of logging on to a workstation. With this feature, administrators can severely restrict the capabilities of users on publicly accessible workstations; for example, they can restrict access to system resources such as hard and floppy disks, and access to specific programs and services. For users who require access to more than one workstation, profiles can be stored on a network server and downloaded to any workstation on the network. Access to a particular profile is granted automatically as a result of the login process, which depends on a user name and password for proper authentication. This feature allows the establishment of a profile for administrators that allows full access, as well as a profile for students that allows highly restricted access.

PROTECTING WORKSTATION COMPONENTS

Depending on the manufacturer, internal workstation components may be easily removed and thus may require special protection. The hard disk, CD, and memory modules are the most vulnerable although damage to any of the internal components could be costly to repair. If tie-downs are used for the workstation, these additional components should be protected as part of the tie-down design. Without tie-downs, the workstation should not allow easy removal of internal components. However, tie-downs must not add additional maintenance costs when units must be serviced.

HARD DISKS

Restricting end-user access to the operating system ensures that original workstation configuration files remain protected. Such restriction may involve third-party desktop security products, policies and profiles, and limited access to the operating system file management system.

On NT systems, NTFS is the recommended file system for ensuring optimum security.

BIOS

Internal hardware configurations on Intel platforms are initially set through the system BIOS and thus should be password-protected to ensure the integrity of the settings. Of concern is Internet access to BIOS "cracking" utilities, which can effectively reveal the BIOS password.

On Intel platforms, the default boot diskette is Drive A. This can be redefined as Drive B to restrict system booting from Drive A thereby preventing the transfer of boot sector viruses. The downside of this change is that technicians can not boot service software from drive A.

COMPUTER VIRUSES

When a virus attacks an animal, it latches on to one of the components that is essential to the reproduction of cells and makes use of that component to facilitate its own reproduction. A computer virus operates in a similar way. A computer virus is a program that attaches itself to or replaces part of another program, and then makes copies of itself whenever the opportunity arises. The virus will accomplish its objective in a way that is not immediately obvious, thereby giving itself the best possible opportunity to continue spreading.

Many extremely destructive viruses were designed specifically to destroy data or cause some other disruption. However, most viruses are "comparatively harmless." They simply wait for an opportunity to reproduce, and spread without causing any major problems. But even these viruses can affect performance. They occupy disk space and memory, and they use CPU processing time. And, of course, these otherwise "harmless" viruses require an investment in time and money to detect and remove. As well, some viruses that were intended to be harmless can cause the random and unpredictable loss of data.

SOURCES

Anyone who can write computer software can create a virus. Kits are available on the Internet that provide instructions and code segments for creating viruses. Some people use these kits to teach themselves the mechanics of viruses. But anyone with a malicious mind sees such resources as an opportunity to wreak havoc. As the existence of these kits suggests, there are many highly skilled programmers throughout the world who are capable of writing viruses without any aids whatsoever.

ACTIONS

Although viruses all behave in a generally similar manner, they can be categorized according to the specific type of site they attack. A file virus attaches itself to an executable file such as a word processing program. Whenever the program is activated, the virus has an opportunity to replicate. In the early days of viruses, data files were not generally infected because data files are not executable, and they present no opportunity for the virus to replicate. However, data files now often carry information such as macro definitions that can be executed by another program. So-called "macro"

viruses have spread widely over the past few years. Companion viruses, a variation of the file virus, spread via a file which runs instead of the file the user intended to run, and then runs the original file so that the user is unaware of their existence.

Every DOS-formatted disk has information on its first sector (the boot sector) that is used to start the computer's operating system. Boot sector viruses alter this information. When a disk with a boot sector virus is used to start the computer, the virus executes its own code, copies itself to the hard drive, and then allows the startup process to continue. Then, whenever a write-enabled floppy disk is used on the infected computer, the virus copies itself to the boot sector of that floppy disk.

Viruses that exhibit features of both file and boot sector viruses are called multipartite viruses. When an infected file is executed, it infects the hard disk and thereby gives itself the ability to infect write-enabled floppy disks.

MECHANISMS OF TRANSMISSION

Boot sector viruses are normally spread when floppy disks are carried from one computer to another. The PC becomes infected when it is rebooted from the infected floppy disk. Quite often this happens by accident when a computer is rebooted with the floppy disk installed and the BIOS is configured to boot preferentially from the floppy drive. Although boot sector viruses are very common, they can not normally be spread across a network. Infected floppy disks can come from numerous sources including home computers, service technicians, or colleagues. In fact, there have been cases involving the unwitting distribution of viruses by reputable software vendors.

File viruses can spread quite quickly across a network. If someone distributes an e-mail message with an attached file that is infected, every recipient of the message can be infected simply by executing the program associated with the attached file. Such viruses can be readily spread by computer bulletin boards or FTP sites, and they can of course also be spread in the same way as boot sector viruses.

Multipartite viruses have the ability to spread via a network because of their attachment to files. Once they have reached their target, they can infect the boot sector of the target's hard drive.

DETECTION AND PREVENTION

Policy

One of the most important tools for fighting the spread of viruses is an appropriate policy for dealing with files from outside sources. Many organizations forbid the use of floppy disks from outside the organization unless the disks have been scanned with virus-checking software. Other components of the policy should include a process for reporting and dealing with virus outbreaks, standards for virus detection and elimination software, and standards for dealing with files arriving from outside sources via modem or network connections.

Commercial Workstation and Server Detection Software

There are many commercial software packages that are intended to detect and destroy computer viruses. These programs use a variety of techniques to accomplish their objectives. They scan files for specific patterns, they monitor running programs in order to detect activities that may be the result of viruses, and they maintain databases that can be used to detect changes in files that may have been introduced by viruses. The creators of such software are always trying to keep up with new viruses that are constantly appearing, so it is important to regularly update any virus detection software.

Since every machine in an organization is vulnerable to viruses, the best way to protect against damage from viruses is to equip each machine with appropriate detection software. For a large organization such as a school jurisdiction, this can be an expensive and time-consuming activity because such software must be regularly updated. Some manufacturers of such products have recognized this problem, and they have developed server-based software that automatically downloads new revisions to every machine on a network.

Gateway Detection Software

Organizations without the resources to equip every computer with antivirus software, and organizations that want an extra degree of protection can make use of products that examine files received from an outside network via e-mail or FTP. These products scan the incoming files and prevent suspicious files from being transmitted onto the network until they can be fully examined. Since viruses can be hidden in a great many ways, such products are not foolproof, but they do have the capability of stopping many of the commonly occurring viruses. Even if a virus spreads to only one or two machines, the amount of effort required to eliminate the virus and recover from its effects can be significant, so anything that can be done to detect a virus at its point of entry is well worthwhile.

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

The student population is curious and resourceful, and sometimes a few students have bad intentions. It is difficult enough to fully control physical access to critical computers. It is almost impossible to fully control physical access to computers that students use regularly.

Therefore, a clear, strongly enforced policy concerning the physical abuse of computers is really the only effective mechanism for protection. Such a policy should specify which activities are allowed and which activities are specifically forbidden as well as which computers are accessible to students and which are not. It should describe how physical access will be controlled and monitored, and specify penalties for contravention.

Many students now have computers at home and need to transport floppy disks from home to school. As a result, the importation of viruses is a major concern. The best way to combat this problem is to equip every computer with antivirus software. Failing that, a number of readily accessible computers should be equipped with up-to-date

antivirus software, and a policy established and enforced regarding the scanning of all outside disks.

Antivirus software on every computer is the best mechanism for dealing with files imported from the Internet or other sources. When such software is not available, it is important to restrict access to outside networks by eliminating dial-up modem connections and providing a single Internet gateway. If at all possible, this gateway should be protected by some sort of virus detection software.

REFERENCES

Virus Myths and Hoaxes

- <<http://www.kumite.com/myths/>>
- <<http://ciac.llnl.gov/ciac/CIACHoaxes.html>>

Viruses and Windows NT

- <<http://www.symantec.com/avcenter/reference/vbnt.html>>

SECTION III

THE LOCAL NETWORK

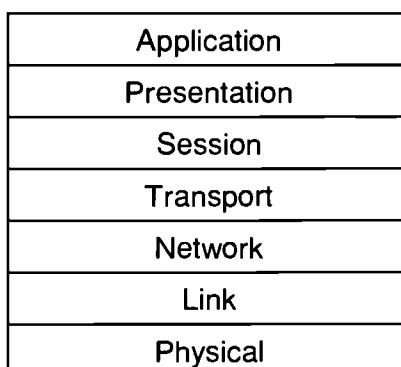
THE BASICS

Just as a locksmith needs to know about the internal workings of a safe in order to ensure that the safe is working properly, a network administrator must understand how information is transmitted throughout the network, and what mechanisms can be used to intercept or interfere with that information. A standard model, referred to as the OSI Reference Model, has been adopted as a means of explaining how the various components of a network interact. The model demonstrates how the network can be viewed as a series of layers that interact with one another. Each layer is responsible for a specific function, and it has a standard interface to the layers above and below it. This model makes it possible to see how information is transmitted from a specific application through the various levels of the data communication process to the actual physical cable which transmits the information, and then back to the application receiving the information.

A well-designed network will ensure that information is transmitted from its source to its destination as effectively as possible, and it will minimize the potential for that information to be damaged or intercepted. Additional measures such as encryption can be employed to further ensure data security.

THE OSI REFERENCE MODEL

Transmitting information from one computer system to another involves a series of interrelated activities. In the early 1980s the International Organization for Standardization developed an architectural model for the process. This model, called the Open Systems Interconnection (OSI) model, represents the entire process as a series of seven layers. Each layer is responsible for a specific component of the process, and each layer communicates with its neighbours using standard mechanisms. The OSI model has become almost universally accepted as the means of describing the data communication process. Even though some manufacturers have developed proprietary protocols that do not exactly match the layers specified in the model, they tend to document their protocols with reference to the model. The model is most often depicted as a stack with seven layers



A detailed description of the functions of each of these layers can be found at:

<<http://www.rad.com/networks/1994/osi/layers.htm>>

The OSI model is similar to the process of sending a letter across the country via a courier service. The writer places the letter in an envelope with the name and address of the recipient. That envelope is then forwarded to the mail room, where it is placed in a courier envelope, and the courier envelope is placed in a bag with all of the envelopes destined for a particular city. The bag is carried to the destination city, and the courier envelope is removed from the bag, delivered to the mail room, and opened. The envelope is then delivered to the recipient. Every placement of the message in another container is analogous to the transfer of information from one level to another in the OSI model.

One other point to notice: each layer deals only with a particular type of container. The courier company knows what to do with mail bags, the mail rooms know what to do with the courier envelopes, and the secretaries know how to deal with the envelopes containing messages.

In this example, the messages are put into different containers when they change levels. An electronic message consists of a string of binary digits, and when a message is transmitted from one level to the next, an extra series of digits called a header may be attached to the beginning of the message. At the receiving end, the headers are removed one by one as the message passes from the physical level to the upper level. Each level deals with the header created by the corresponding level on the transmitting side. The headers contain information that allows the message to find its way from the source to the destination.

Most network security problems are associated with the three lowest levels.

The physical layer defines the actual physical mechanism by which the information is transmitted. It includes specifications for such things as voltage levels, timing, physical connectors, and media.

The link layer is responsible for ensuring that data is transmitted reliably across the physical link. It deals with topics such as physical addresses, network topology, error notification, and flow control.

The network layer is responsible for establishing connectivity and selecting appropriate paths between two sub networks that may be geographically and topologically separated.

PROTOCOLS

The following discussion of the more popular networking protocols is intended to provide background information for future discussions of potential network security problems. Each of the protocols is quite complex. These brief descriptions give only a very basic overview of the features and capabilities of the various protocols. Further information concerning each protocol can be found in the references specified at the end of this section.

ETHERNET

Over the years, Ethernet has become the most popular technology for departmental networks. Each Ethernet-equipped computer on the network operates independently of the others—there is no central controller. All attached stations are connected to a shared media system, and signals are broadcast over the medium to every attached station. A station wishing to transmit information listens for traffic on the medium, and waits until the medium is idle before sending its message. If two stations attempt to transmit a message at the same time, they each detect that a collision has occurred, and they both wait for a random amount of time before trying to retransmit.

The Ethernet protocol can function using several different physical media. These include thin and thick co-axial cable, twisted pair wires, fibre optic cable, and wireless networks; and there are devices that facilitate transmission from one medium to another. Each medium has its own requirements and limitations. When considered with respect to the OSI model, the physical medium would be a component of the physical layer (Layer 1). The rules for accessing the medium, together with the process for building and interpreting the packets of data that travel across the network can be considered as components of the link layer (Layer 2) of the OSI model.

Every manufacturer who builds interfaces that connect to Ethernet networks builds into the interface a unique hardware address that is different from every other address ever issued anywhere in the world. This address is referred to as the Medium Access Control (MAC) address. When the Ethernet protocol receives a message from Layer 3 of the OSI stack, it adds to the message a header that includes the source and destination MAC addresses for the message. The receiving station strips off that header before passing the message on to the next layer.

TOKEN RING

Token Ring technology was developed by IBM in the 1970s as an alternative to Ethernet. Like Ethernet, Token Ring functions at Levels 1 and 2 of the OSI model. But unlike Ethernet, the protocol was designed so that collisions resulting from two stations attempting to transmit at the same time are not possible. In a Token Ring network, stations can be considered to be logically connected in a ring so that each station has two neighbours. A small frame (a string of bits) called a token is passed from one station to the next around the ring, and a station has the right to transmit an information packet only when it has possession of the token. The information frame is passed from one station to the next along the ring until it reaches its destination where it is copied. It then continues around the ring to the sending station, where it is removed. Also included with the Token Ring technology is a priority system that allows more important stations to take precedence over less important ones.

The actual physical implementation of a Token Ring network involves a device called a multistation access unit (MSAU). Each workstation is connected via a direct connection (usually twisted pair) to a port on a MSAU, and each MSAU has an input and an output port by which it is connected to adjacent MSAUs in a ring configuration.

Token Ring networks employ various fault management mechanisms that can isolate a faulty workstation from the ring and deal with network faults.

The popularity of Token Ring has declined in recent years, so there are relatively few new installations, but the technology is still widely used.

TCP/IP

The TCP/IP protocols, upon which the Internet and most local networks are based, operate at Levels 3 and 4 of the OSI model. At Level 3 is the Internet Protocol (IP) which deals with routing of packets, fragmentation and re-assembly of packets, and error reporting. The Transmission Control Protocol (TCP) resides at Level 4 of the OSI model. It provides service to upper layer protocols in such a way that transmission and reception can occur simultaneously. It provides error detection and correction mechanisms, and it can support simultaneous conversations with upper layer applications.

Just as every device on a network has a unique MAC address, it also has a unique IP address. The difference is that the MAC address is inextricably associated with the hardware components of the Ethernet interface whereas the IP address is assigned via software and can be modified at will. The IP protocol communicates using IP addresses, whereas the Ethernet protocol communicates using MAC addresses. The Address Resolution Protocol (ARP) helps to provide an interface between the IP and Ethernet protocols. When the Ethernet protocol needs to know the MAC address associated with a particular IP address, it broadcasts an ARP request to all stations on the network. The station with the requested IP address responds with a message that includes its MAC address. Subsequent conversations between the sending and receiving stations can then take place using the returned MAC address. Most devices maintain a table, often called the ARP table, that keeps track of the results of ARP requests so that once a station's MAC address is determined, further ARP requests are not needed.

NETWARE

NetWare is a network operating system that was created by Novell, Inc. in the early 1980s. It is particularly suited to small workgroup-type networks that share resources such as file and print services, although it also has been employed in much larger installations. It operates at the upper five layers of the OSI model. With appropriate drivers, it can interact with almost any media-access protocol. The Level 3 protocol associated with NetWare is called Internet Packet Exchange (IPX). IPX is capable of routing packets from a source network, through a series of interconnecting networks, to its destination. The most commonly used NetWare transport protocol is called Sequenced Packet Exchange (SPX). Like TCP, it provides error detection and correction and connection services.

NetWare's current implementation of NWIP is an IP-encapsulated IPX packet required for communication between NetWare clients and servers in an IP-based environment. With the announced release of NetWare 5.0, native IP will be available.

APPLETALK

The design of AppleTalk allows it to function with any link-layer implementation. Manufacturers of network equipment have taken advantage of that feature by developing AppleTalk implementations for all of the popular network media. The protocol was designed to minimize administrative overhead; it accomplishes automatically many of the tasks that require manual intervention in other protocols. This requirement for minimal administration has resulted in a need for more communication among network devices, so AppleTalk is considered to place more demands on network bandwidth than most other popular protocols. AppleTalk addresses are assigned automatically, and messages can be automatically routed from one network to another in a series of interconnected networks.

NETWORK SEGMENTATION

The simplest device for connecting various components of an Ethernet network is a multiport transceiver, or hub. This device has multiple interfaces, each of which can be connected to a separate segment of the network.

As mentioned earlier, the Ethernet protocol uses broadcasts to transmit information. This means that, as more and more stations are added to a network, more and more traffic is generated. Since the network has a finite capacity for carrying information, there is a limit to the number of stations that can be attached to a particular network. To overcome this problem, a network is divided into two or more sub networks when it grows to a certain size. Stations on each sub network broadcast only to devices on the same sub network, and communication with other sub networks takes place via devices which transmit only traffic destined for those sub networks.

The simplest of these devices is a bridge. It possesses a separate Ethernet interface for each sub network to which it is connected, and it inspects each Layer 2 Ethernet packet and decides to either forward or block the packet on the basis of its destination MAC address. When the bridge is first turned on, it listens to traffic on the sub networks to which it is connected, and keeps track of which MAC addresses are on which sub network. It uses this information to decide which sub network to forward a packet to.

A more complex device that functions much like an Ethernet bridge is called a Layer 2 switch. To the casual observer, it looks like a hub, but internally it is functioning as a bridge so that each port on the switch is attached to a separate sub network. Switches have the capability of establishing virtual LANs (VLANs). These are sub networks that may involve stations on different physical network segments. Generally, VLANs are used in situations where there are many sub networks, each with only a few stations. The switch forwards traffic to all stations on the VLAN, no matter where they are located, but the traffic is not forwarded to sub networks that do not have stations associated with the VLAN. There is another major benefit of switches: the only traffic on an individual segment is the traffic associated with devices on that segment. So a switch can be useful for breaking up a congested network into a number of less-congested segments. As an extreme example, if only a single device were attached to each port of the switch, each device would be isolated on its own network segment and be able to use the entire traffic-carrying capacity of that segment.

Because a bridge or a switch has a finite capacity for remembering MAC addresses, the use of these devices can not be expanded to an extremely large network such as the Internet that involves millions of MAC addresses. To communicate effectively in such circumstances, a different approach called routing has been developed.

An IP address is thirty-two bits long. The protocol specifies that this address must consist of two parts; one part represents the address of a particular sub network and the remainder represents the address of a particular station on the sub network. The address associated with the sub network is called the network address, and the address of the station is called the host address.

Routing makes use of network addresses. Devices called routers have multiple network interfaces, and each interface is associated with a particular network and therefore a particular network address. Routers communicate with one another, so they are aware of the best path to particular sub networks in their vicinity. Each router also has a destination referred to as a default gateway where it sends packets that do not have a specific destination.

The Internet can be viewed as a tree. The tips of the roots are individual workstations, and the trunk represents a group of routers, referred to as core routers. The routers are traffic directors located wherever a root branches. Each router knows how to get to networks below it, and the router at the lowest level knows how to contact every workstation on the networks to which it is attached. If a router can not locate the destination for a particular packet, it uses its default gateway to pass the packet up to the next higher branch. This continues until the packet reaches a router that knows where to send it.

Routers make use of IP (Level 3) addresses, whereas bridges work with MAC (Level 2) addresses. Because routers must communicate with one another, and because they have to disassemble the packets in order to inspect them, they tend to be slower than bridges. A new development is a device called a Level 3 switch that performs the same functions as a router but is much faster because of differences in internal design. Currently, Level 3 switches from vendors may not be standards based. Careful research is required before selecting products.

NETWORK MANAGEMENT

SNMP

Simple Network Management Protocol (SNMP) was developed so that network administrators could remotely monitor and control network equipment such as bridges, routers and switches. The protocol makes use of TCP/IP to send and receive messages that contain instructions for network devices or information about status and network traffic. Network management systems employ software on a central system referred to as a Manager to interact with agent software on the remote devices. The agent software interfaces with components on the remote device that are referred to as objects. These objects might be the actual hardware, the configuration parameters for

the device, or performance statistics. The objects are arranged in a structured manner in a database referred to as a Management Information Base (MIB).

Most MIBs are designed to work with a particular piece of network equipment, but the Remote Network Monitoring (RMON) MIB was designed to allow management of the entire network. RMON agents can be stand-alone devices, or they may be incorporated into a network device. RMON allows network administrators to obtain detailed information about the quantity and type of traffic being carried by a particular network segment. Careful research is required before selecting products as devices may have different levels of RMON support.

The most commonly used version of SNMP (Version 1) is fairly limited with respect to security capabilities. A simple clear-text password called the "community string" is used to obtain access to the agent, and then the agent decides whether to accept a particular instruction on the basis of the sender's IP address. Since most network devices are shipped with the default community string made public, one of the first steps in installing any new device is to reset the password.

SNMP Version 2, which is incompatible with Version 1, was designed to overcome the security weaknesses of Version 1, but it has not been widely implemented.

ICMP

The Internet Control Message Protocol (ICMP) was designed to report failures in routing back to the source of the packet being routed. The reporting is accomplished by special messages such as:

- echo and reply messages that test node reachability across an Internetwork,
- redirect messages that cause packets to be routed efficiently,
- time exceeded messages that inform sources that a packet has not reached its destination in the time allowed for it to do so,
- router advertisement and router solicitation messages that allow routers to know about other routers on the network.

One of the most useful tools for any network administrator, a program called "ping," employs ICMP to test whether information can be transmitted to a particular network address. Another useful program, "traceroute," uses ICMP to report the path to any location on the network.

Although such tools are extremely valuable resources for the network administrator, they also can be used by unauthorized individuals to collect information that may be useful for penetrating the network.

NETWORK SNIFFING

Network sniffing is attaching a device to the network to monitor all traffic. Since all of the popular protocols depend on broadcast information, there is no easy way—short of encryption—to prevent data from being intercepted.

A switch prevents data sniffing since it will forward data only to the intended port and not across an entire network. If workstations are dedicated per switch port, these independent workstations are relatively secure, but some switches have sniffing ports. Therefore, physical access to the switches is required in order to connect to the sniffer port.

TOOLS

There are two ways to accomplish network sniffing. A separate device can be attached to a network, or special software can be loaded on an otherwise legitimate workstation. Stand-alone sniffers are marketed as network diagnostic tools, and they have sophisticated capabilities for capturing and analyzing network traffic. These devices are not normally a threat because they are quite expensive. The exception is sniffing networks where a financial gain could be obtained. However, software can be downloaded from the Internet that allows any personal computer to be used as a sniffer. So any portable computer, or even a computer with a legitimate network connection, can be used to collect information from the network.

POTENTIAL VULNERABILITIES

Since network sniffing is one of the most powerful tools for extracting information from a network, it is very important to prevent it. Good physical security of wiring closets and network equipment can reduce the opportunity for connection of a stand-alone sniffer to the network. This potential can be further reduced by installing hubs and switches that will react to the connection of an unauthorized device by disabling the port to which it is connected. Since a sniffer can be attached to any functional network access point, disabling any connection points that are not regularly used is an important security precaution.

It is more difficult to prevent someone from loading sniffer software onto a legitimate network workstation. The potential for this type of sniffing can be reduced through the use of profiles that restrict administrative access and disable the execution of programs from floppy disks.

Network segmentation also is a very useful mechanism for keeping information from falling into the wrong hands. The segment of a network that is attached to a particular port on a switch sees only traffic to and from the devices on that segment. If workstations or network ports that may be potentially compromised are isolated on their own segment, a sniffer attached to that segment will not see traffic flowing on other segments. For example, if a student lab is on its own network segment, traffic flowing from a teacher's office to an administrative system can not be viewed by a sniffer in the lab.

Network segmentation makes troubleshooting more difficult because the network administrator can not remotely view traffic on an isolated segment to detect problems. Therefore switch manufacturers are now incorporating a mechanism that allows an administrator to remotely "tap into" any segment served by the switch. Obviously, it is very important to ensure that the passwords and procedures for such a process are kept secure.

DATA ENCRYPTION

Encryption is one of the most politically sensitive topics in the field of data communication. Some governments are concerned that strong encryption techniques falling into the wrong hands could threaten the stability of the countries they govern. In fact, the United States government has placed restrictions on exporting certain encryption technologies. Despite the potential problems, encryption is one of the most useful tools for network security.

Encryption involves a concept referred to as a "key." A key is a series of binary digits used together with an encryption algorithm to encode a message. So long as the person who receives the message has access to the appropriate key, s/he can apply a decoding algorithm and recover the content of the encrypted message. Keys also can be used to sign messages so that the recipient can be sure of the source. Since keys are so important, they must be guarded carefully. The process of issuing and managing keys has become a major industry.

SPECIFIC INFORMATION

Since data encryption involves extra software, overhead, and cost, this technique is normally used only where it is deemed to be cost-effective. One of the most important types of information that could be collected from a network is a password. For that reason, some security policies insist that free text passwords must not be transmitted via the network, even though other traffic is transmitted openly.

Another application is the World Wide Web. There are browsers that can encrypt traffic as required in order to safeguard the transmission of such things as passwords or credit card numbers.

ALL INFORMATION

On a virtual private network, all traffic between specific sites on a network is encrypted. Such networks allow a public facility such as the Internet to be used as a medium for transporting confidential information. At the access points to the network, routers or firewalls encrypt and decrypt information.

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

NETWORK TOPOLOGY

Data communication networks for school jurisdictions perform two distinct functions: student-related and administrative tasks. Since the greatest potential for problems is associated with students, it is highly advisable to physically separate student and administrative networks as completely as possible. The appropriate use of switched network segments and VLANS can be very useful in this respect. In fact, because the cost of switched ports has been reduced considerably in recent years, it generally is advisable to use switches instead of hubs in most circumstances. A firewall between student and administrative networks also can provide added security. Where such

separation is not practical, encryption and virtual private networks can be useful tools for protecting confidential information.

IP ADDRESS ASSIGNMENT

School jurisdictions tend to need large numbers of IP addresses because they have a large number of devices that require addresses and there are geographical considerations (their networks are distributed over many sites). IP addresses are now in short supply. It is generally not possible to acquire enough addresses to allow an efficient network implementation.

The most effective solution to this problem is to use a technique known as IP address translation, described in RFC 1918. With this technique, the internal network uses non-routable IP addresses, and traffic for locations outside the network uses a router or firewall that translates the internal address to an assigned, routable address. Three specific blocks of non-routable addresses have been defined, and these addresses provide more than enough capacity for even the largest school jurisdiction. The addresses that can be used for this purpose are:

- 10.0.0.0–10.254.254.254
- 172.16.0.0–172.31.254.254
- 192.168.0.0–192.168.254.254

Implementing dynamic IP addressing through DHCP can reduce the management and record keeping of workstation addresses. Still, assigning static addresses to specific stations may be useful for security access permissions; i.e., internal firewall rules.

REFERENCES

General References

[<http://www.dc.net/ilazar/default.htm>](http://www.dc.net/ilazar/default.htm)

Ethernet

[<http://wwwhost.ots.utexas.edu/ethernet>](http://wwwhost.ots.utexas.edu/ethernet)

Network Management

[<http://www.cis.ohio-state.edu/hypertext/faq/usenet/snmp-faq/part1/faq.html>](http://www.cis.ohio-state.edu/hypertext/faq/usenet/snmp-faq/part1/faq.html)

Cryptography

[<http://www.csua.berkeley.edu/cypherpunks/Home.html>](http://www.csua.berkeley.edu/cypherpunks/Home.html)

SECTION IV

SERVERS

THE BASICS

No matter how well a server is protected from access by a network, it is highly vulnerable if it is not physically secure. Also, because administrative passwords allow access to the very heart of the facility, it is important to establish appropriate password management procedures. Where file systems are shared among users, the permissions to access such systems must be properly established and managed.

Each server operating system has its own particular security weakness. Therefore, it is important for system administrators to have a good understanding of the operating system as a whole and of the potential security holes and possible solutions.

PHYSICAL SECURITY

Physical security is important for the protection of workstations, but it is essential for the protection of servers. Access to a server allows access to the privileged information that is stored on the server, and it also provides a wide open door to the rest of the network to which the server is attached. All servers should be placed in secure locations.

Because servers often act as repositories for essential information, and because they often control the operation of the entire network, they need to be protected from dangers in addition to those posed by malicious humans. Essential servers should be equipped with uninterruptable power supplies that are capable of providing protection against power surges. Wherever possible, they should be placed in an environment where the climate is controlled and monitored, and they should be protected from the effects of potential water damage resulting from broken pipes or floods.

ADMINISTRATIVE PASSWORDS

All of the precautions recommended for password management on workstations are even more important for servers. Only a limited number of people should have access to administrative passwords; those passwords should be long enough and complex enough that they can not be guessed or cracked using standard cracking programs, and they should be changed regularly. Passwords should never be written down, transmitted via e-mail, or transmitted verbally via cellular telephone.

When selecting a password, staff should consider the following guidelines:

- Choose a password that is at least seven characters long.
- Use a mixture of upper and lower case letters, numbers, and punctuation symbols.
- Do not use personal information such as telephone numbers, license plate numbers or names of friends, relatives, or pets.
- Do not use words or names that can be found in a standard dictionary.

- Avoid using passwords that can be easily observed as you type them; e.g., aaaaaa or 1234567.
- Never use your user ID as your password.
- Do not re-use passwords that were in recent use.
- Use passwords that are easy to remember. Some people use the first letter of each word of an easy-to-remember phrase; e.g., "My dog has lots of fleas!" would result in the password "Mdhlol!"

Another important password management practice is to change passwords regularly. Most authorities recommend changing administrative passwords at least once a month. In addition, it is essential to change passwords immediately when any administrative staff leave the organization or when there is the least suspicion that a password may have fallen into the wrong hands.

SHARED FILE SYSTEMS

As a general rule, access to a shared file system should be granted only to users who need access. Also, those users should have permissions that provide only the rights they need to do their job. Every network operating system has its own potential weaknesses with respect to shared file systems. The following pages provide general descriptions of the most common operating systems and security precautions associated with those systems.

ELECTRICAL POWER

As a general rule, any server that supports a production application should be equipped with some sort of protection against power failures and fluctuations. Sudden removal of power from a server can cause loss of data, and power surges or dips can result in data loss as well as physical damage to the server. Uninterruptible power supplies (UPS) are available in a wide range of sizes, ranging from book-size units intended for providing a few minutes of protection to a workstation all the way up to rooms full of equipment and batteries designed to protect an entire building. This is the first question in establishing a computer installation: do we provide the entire computer room with uninterrupted power, or do we provide separate protection for each server in the room?

Generally, a UPS for an individual server is intended to provide enough power to allow the server to survive an outage lasting a few minutes, and to allow a graceful shutdown if the outage is longer than the capacity of the UPS. This type of approach is relatively inexpensive and is most appropriate for non-essential servers and in situations where the local power supply is reliable. For situations where failure of a server is not acceptable or the local power is unreliable, a more elaborate installation is required. This type of installation usually incorporates a battery-powered UPS backed up by a diesel generator that can continue to provide power so long as it has fuel. This is the type of installation that would normally serve an entire computer room rather than an individual server.

To select a UPS to protect a server, determine how much power is required and the maximum length of time for which protection is needed. The server manufacturer can provide information on power consumption, and the UPS vendor can then use this

information to specify the most appropriate unit. Another important requirement for a UPS is a means for communicating with the server so that the server can initiate a shutdown before the UPS batteries are completely discharged. For situations where the server is not readily accessible, a useful feature is the capability to remotely monitor and/or manage the UPS.

A major component of a UPS is the battery. Since batteries tend to fail over time, it is important to regularly check the UPS according to the manufacturer's instructions to ensure that it is functional and to replace batteries according to the manufacturer's recommendation. The worst possible time to detect a bad battery is during an actual power failure.

RAID

Disk drives are one of the major sources of system failures because they are one of the few components of a modern computing system with moving parts. To address disk problems for installations that require minimal down-time, manufacturers have developed products based on a technology referred to as RAID (Redundant Array of Independent Disks). RAID spreads data over two or more disks in such a way that a single disk failure will not cause a complete system failure. Several different RAID implementations have been proposed, but only two—RAID 1 and RAID 5—are in common use. RAID 1, also referred to as disk mirroring, writes data onto two separate disks so that two copies of the data are always available. This means that the actual per byte cost of storage is doubled with RAID 1. RAID 5 spreads the data across a group of disks, and it incorporates some extra information that allows the data to be recovered if any one disk in the group fails. The cost of storage for RAID 5 is less than for RAID 1 because only enough extra disk space to hold the error-correcting information is necessary. However, because the price of an individual disk is now relatively low, RAID 1 has become the most popular approach. The actual flow of information from the processor to the disks is handled by the RAID controller, so the impact on system speed is not usually significant.

DATA BACKUP

Almost everyone who uses a computer regularly has lost a critical file. Such losses occur through operator error, hardware or software failures, or unauthorized access to the computer. No matter how the loss occurs, the first question is, "Did you do a backup?"

Since servers act as the major repositories of information for most organizations, server backups are absolutely essential. Depending on the numbers and sizes of the servers, backup processes can take a number of forms. For small servers, a regular manual backup may be sufficient, but large installations usually employ special backup software that automates the process as much as possible.

Various physical media can be employed for data backups, including a variety of magnetic tape and disk formats, and even writeable CDs. Most production environments use some sort of cartridge tape system, and the normal approach is to back up everything on the system at some regular interval, and then back up on a daily basis only information that has changed since the previous backup. The backup interval

and the exact scheme for doing the incremental (daily) backups depends on the value and reproducibility of the data, the availability of a backup device, and the maximum time desired for restoring the data.

No matter which backup process is chosen, one major issue should not be overlooked. It is absolutely vital to ensure that the backup process is actually doing what is expected of it. The only way to be sure is to periodically test the process. An actual restore should be completed annually.

UNIX

The first version of UNIX was created in the early 1970s by two researchers at Bell Labs. Over the years, numerous variations have evolved as different organizations have added their own features and capabilities. Today, each manufacturer of computer systems that are designed for UNIX offers its own version of the operating system. Although all of these versions function in essentially the same manner, each has different features and different vulnerabilities. Therefore, any discussion of UNIX security must refer to specific versions of UNIX.

VARIATIONS OF UNIX

The UNIX family tree has two major branches: versions created by the University of California at Berkeley and by AT&T, who attempted to develop a commercial version. The Berkeley version is normally referred to as BSD, and the AT&T version as System V (where V is the Roman numeral five). Most current versions of UNIX can trace their ancestry to one of these two sources.

Following is a list of the most widely used versions of UNIX.

Sun Microsystem's Solaris 2.X is an AT&T UNIX with many extensions. SunOS, an older operating system from Sun, is a BSD-based system, also with many extensions.

DEC is the only major manufacturer to have adopted the OSF/1 standard, which is in turn based on the Mach operating system developed at Carnegie-Mellon University. MACH is derived from BSD, so there are many administrative similarities between Digital UNIX and other BSD implementations.

Older DEC machines run DEC's original version of UNIX, called Ultrix, which is very closely linked to BSD UNIX.

IBM's version of UNIX, called AIX, exhibits characteristics of both System V and BSD as well as many unique features that have been developed by IBM.

IRIX, marketed by Silicon Graphics, was originally similar to BSD UNIX, but current releases tend to have more and more System V features.

Hewlett-Packard's system (known as HP-UX) is derived from System V, but it also has some unique features.

One of the early commercially successful versions of UNIX for Intel-based personal computers was developed by SCO, the Santa Cruz Operation. This operating system is based on an old version of ATT UNIX, System III.2. It has been modified extensively.

Linus Torvalds of Helsinki developed a system called Linux that is freely available and runs on Intel-based PCs. Many developers are now creating extensions for Linux. Its internals and system administration are similar to BSD versions, but its programming interface is more like AT&T UNIX.

NetBSD, 386BSD, and FreeBSD are other PC UNIX implementations based on Berkeley UNIX.

COMMON SECURITY PRECAUTIONS

Although each version of UNIX has its own idiosyncrasies and weaknesses, every UNIX systems administrator must take certain precautions, as described below.

General Security Policies/Procedures

- Install vendor-supplied security patches promptly.
- Monitor account security regularly.
- Make sure someone receives security bulletins and notes from the systems vendor and from organizations like CIAC and CERT.
- Monitor system log files regularly.
- Set policies for users and make sure they understand and abide by the policies.

Account Security

- Tell users what constitutes an effective password and use either a proactive password checker or a password cracking program to verify that passwords are secure.
- If desired, and your system supports it, implement password aging.
- If your system supports it, implement a shadow password file. Many systems offer the capability of establishing a shadow password file; i.e., a file in a secure location that contains encrypted user passwords. This prevents an intruder from capturing the password file and using an off-line cracking program to obtain passwords.
- Monitor account usage, and disable accounts that are not being used.
- Minimize the number of accounts on servers and "critical" hosts.
- Minimize the number of users with "super user" privileges.

Network Security

- Check the CERT Archives for advisories regarding known problems with your version of UNIX.

- Ensure that services such as sendmail and FTP have the latest security patches and are properly configured.
- Disable all unnecessary services; e.g., TFTP or finger may not be required on the system.
- If supported, enable logging of successful and failed network connections. Better yet, use the TCP wrapper or xinetd program for logging and to allow only specified hosts to obtain access to network services.

Physical Security

- If possible, install a version of the PROM monitor that either does not provide (or at least password protects) the commands to examine and change memory contents.
- If possible, ensure that workstations can not be taken into single-user mode without providing the "root" (or a PROM monitor "hardware") password.

File System Security

- Set appropriate file permissions on all files.
- Ensure that default file permissions are set appropriately for each user.
- Do not allow set-user-ID or set-group-ID shell scripts on the system.
- Check all "nonstandard" set-user-ID and set-group-ID programs for security.
- Implement and use a comprehensive backup scheme.

The items specified above represent some of the more important steps that should be taken. A more detailed list can be found at:

<<http://stimpy.cac.washington.edu/~dittrich/R870/security-checklist.html>>

OTHER SERVER OPERATING SYSTEMS

Windows NT	See Appendix II
Novell	See Appendix III
Apple	See Appendix IV

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

In a school jurisdiction, the greatest computer security threats come from within. Since servers represent the heart of any network, it is extremely important to protect servers from potential abuse from within the network. Physical security and password protection are even more important in school jurisdictions than they might be in other enterprises. Ensure that servers are regularly monitored for signs of abuse and that system administrators keep up to date on new security vulnerabilities. A stringent backup process should be in place.

REFERENCES

Windows NT Security

<<http://www.iss.net/vd/sitesn.html>>
<<http://www.microsoft.com/security/ntprod.htm>>

Uninterruptible Power Supply Buyers Guide

<<http://www8.zdnet.com/pcmag/features/ups/upstest.htm>>

UNIX Security

<ftp://ftp.auscert.org.au/pub/auscert/papers/unix_security_checklist>
<<http://www.alw.nih.gov/Security/Docs/network-security.html>>
<ftp://info.cert.org/pub/tech_tips/UNIX_configuration_guidelines>

SECTION V

REMOTE ACCESS

THE BASICS

Most network administrators regularly face the question, "How can I connect to the network from my computer at home?" Remote access can be accomplished in a variety of ways but each method poses somewhat different security challenges. No matter which technology is chosen, it is important to implement a secure authentication process and a comprehensive logging mechanism.

TECHNOLOGIES

DIAL-UP

The most common method for remotely connecting to a computer network involves the telephone. As with many data communication technologies, an acronym has appeared: POTS, for "Plain Old Telephone System."

The critical component at the client side of a POTS connection is a modem. External modems have a connection for a telephone line and another serial connection that can be attached to a terminal or a serial port on a personal computer. Internal modems are mounted inside a personal computer.

Over the years, modem speeds have increased dramatically. Currently, the accepted standard is 33.6 Kbps. Modems that operate at 56 Kbps also are available now, and manufacturers are moving toward a standard that will allow modems from different manufacturers to be interchangeable. There is one restriction on these devices: they will communicate only with special-purpose terminal equipment that is serviced by a high-speed digital connection from the telephone exchange.

The simplest way of establishing a remote connection is to connect a modem to a device on the network that has a serial port, and then allow external users to establish a dial-up connection with the modem. To expand this facility, more modems and more telephone lines can be added. A more sophisticated and easily managed approach involves using special-purpose remote access servers that include banks of rack-mounted modems and a means of connecting directly to the internal network.

ISDN

ISDN (Integrated Services Digital Network) has become available in most areas over the past few years. ISDN involves the transmission of digital signals over the type of wiring normally used for POTS service. Normally, an end-user of the ISDN service would order something called a Basic Rate Interface (BRI). This service provides two separate channels (referred to as B channels), each with a capacity of 64 Kbps. Depending on which type of equipment is connected to the circuit, ISDN can be used for voice, video, or data transmission. If necessary, multiple BRI circuits can be combined to allow higher data rates than those that can be achieved from a single circuit.

ISDN is a dial-up service, similar to POTS. Each channel of a BRI has a separate telephone number, so a site with ISDN capability can connect to any other site with similar capability. Because ISDN is particularly suited to data transmission, a common practice involves establishing a connection from a remote site to a central location, leaving the connection established, and then using it as though it were a permanent hard-wired connection. Another common use for ISDN is as a backup for a faster wide area network connection. If the faster connection fails, the ISDN connection is established. The time required for an ISDN connection is usually only a second or two, so ISDN also can be used to add extra capacity on demand to a wide area link.

Several different types of ISDN equipment can be used for data communication. If a remote site has only a single computer, it is possible to acquire an internal or external ISDN modem that functions much like a POTS modem. There also are ISDN modems available that allow the use of one B channel for voice or fax traffic while the other one is being used for data. If the remote site has more than one computer, it is possible to use a small router with an ISDN port and a port that can be connected to a local area network. At the central site, depending on requirements, it is possible to use individual ISDN routers, or more complex devices that connect to a Primary Rate Interface (PRI) from the telephone company. A PRI provides as many as twenty-three B channels over a single circuit.

CABLE MODEMS

Recently, cable television companies have been offering Internet service via special modems attached to their TV cables. This service has the potential of being very fast, depending on the number of subscribers using it in a particular area, so it has great appeal for remote users who have a need to connect to a network.

There are two approaches to establishing this connection. The first involves using the Internet as the means of communicating between the remote and central sites. Security precautions for an approach such as this are identical to the precautions needed for Internet security, with one added concern. If the information being transmitted between the two sites is confidential, some sort of encryption should be implemented to safeguard the information. The second approach requires the co-operation of the cable company, and it may not be available in all areas. Some cable companies have the capability of creating a virtual private network between two cable modems on their network. Information is encrypted by the cable modem when it enters the cable network, and it is decrypted by the other cable modem when it leaves the cable network.

xDSL

In some locations, users who require high-speed remote access can implement a service referred to as xDSL. DSL refers to a service called digital subscriber line, and the x in the acronym can be replaced by several letters, each referring to a different implementation of the service. xDSL is a point-to-point service; the connection is in essence a direct connection between two sites. The only way an xDSL connection can be rerouted is through direct action by the telephone company. Consequently, this technology represents much less of a security problem than a POTS or ISDN

connection, which can be reached by anyone who dials the appropriate telephone number.

Telephone companies are offering an Internet connectivity service using ADSL which is one of the xDSL implementations. This type of connection requires the same security precautions that apply to any Internet connection.

REMOTE ACCESS SERVERS

DIRECT ACCESS

The simplest means of providing remote access to a network is to install a modem on a server or workstation on the network. The remote user dials the modem and establishes connectivity with the device on the network. Depending on the capabilities of the network device, communication may be restricted only to that device, or the user may be able to interact with other devices on the network. The device to which the modem is attached is responsible for ensuring security.

TERMINAL SERVERS

A terminal server is a stand-alone device intended to allow serial devices such as terminals or modems to communicate with the network. Terminal servers normally are responsible for activities such as user authentication and logging. They may accomplish these tasks without any interaction with other devices, or they may communicate with other servers on the network to obtain information necessary for user authentication and to log information associated with connection attempts. A small remote access facility can be established with a terminal server and a bank of individual modems. More elaborate installations involve rack-mounted modems combined with one or more terminal servers and management software.

DSP-BASED ACCESS SERVERS

The latest servers for remote access make use of a technology called digital signal processing (DSP), where the functions of the modems and the terminal server are combined into a single device. Normally DSP devices are attached to the telephone system by means of one or more PRI circuits. The function normally performed by the modems is handled by a digital signal processor that adapts automatically to the type of information it receives from the PRI. A DSP is capable of interacting with POTS-type modems and ISDN modems or routers. It also can provide full-speed connectivity for 56-Kbps modems. Unlike older hardware-based modems, the DSP's functionality can be upgraded via software as new standards are developed.

xDSL

Sites that need to provide xDSL service to a large number of locations can buy rack-mounted xDSL modems.

PROTOCOLS

SERIAL

Anyone who has used a remote access facility for more than five years is familiar with serial connectivity. A terminal or a personal computer with terminal emulation software at the remote site communicates directly with a computer or other device at the central site using the same RS-232 protocol that would be used by a terminal that is directly connected at the central site. The modem at the central site is often connected directly to a serial port on the central computer, or it may be connected to a terminal server which communicates with the central computer using a more sophisticated protocol such as Telnet. Authentication and logging of such connections is performed by the device to which the modem is connected.

PPP

Point to Point Protocol (PPP) allows a remote user to communicate with a central network via a modem or ISDN connection as though the remote user were directly attached to the central network with an Ethernet connection. Software on a device at the remote site interacts with similar software at the central site. When the connection is initially established, a negotiation process establishes various network parameters such as addresses, servers, and block size. The negotiation process also can handle two different types of authentication protocols, so that authentication becomes more transparent to the user.

PROPRIETARY

Numerous proprietary remote access protocols have been developed by manufacturers who have a need for their own unique way of dealing with remote access. For example, remote control programs that allow a remote user to take complete control of a centrally-located PC use protocols designed specifically for that purpose. The security features of these proprietary protocols should be reviewed very carefully before they are allowed with any device at the central site that is connected to the network.

AUTHENTICATION AND AUTHORIZATION

Authentication involves determining that external users are who they say they are. Authorization is the mechanism used to allow access to particular resources on the network.

MECHANISMS

Simple Passwords

The most common authentication method is asking the remote user for a user ID and password. The major fault with this approach is associated with the management of the IDs and passwords. Since users tend to forget passwords, they choose ones that are easy to remember (and therefore easy to guess). They also tend to write passwords on

a piece of paper and leave them near their computers where they can be seen by others.

Hackers use a technique referred to as "social engineering" to persuade users to tell what their passwords are. For example, a hacker calls an unsuspecting user and states that s/he is with the Information Systems Department, and is doing system maintenance that requires knowledge of the user's password. Many users will provide that information without hesitation.

An appropriate security policy can minimize the effects of most of these practices, but the use of simple passwords for authentication can never be considered a secure process.

One-time Passwords

Several manufacturers have developed technology that generates passwords that are useful only for a few minutes. Every time a user requires access to the system s/he makes use of a hand-held device or a piece of software to acquire a single-use password. Such approaches are reasonably secure, but they also are somewhat expensive, because each user must be equipped with a device for generating the passwords and a central facility for managing the system must be purchased.

Dial-back

Many terminal servers and computers equipped with remote access software have the capability of calling back a user who requests a remote connection. This approach allows connections only with devices at telephone numbers known to the central site. This dial-back approach can be a useful addition to a simple password authentication system, but it can be circumvented by someone with sufficient knowledge. Also, this approach is not useful for users who travel and need to access the central network from an unknown location.

CLI

Many people now have a device that displays the name and number of the person calling them by phone. The technology that facilitates this service is referred to as Calling Line Identification (CLI). Remote access servers can use CLI to prevent connections from unauthorized locations. This approach is technically more difficult to defeat than dial-back, but it can probably be circumvented by someone who has the necessary knowledge and equipment. It also suffers from the same disadvantage as dial-back with respect to users who travel.

PAP

PPP can use Password Authentication Protocol (PAP) for user authentication. PAP works in the same way as the simple password approach discussed earlier, but the prompts and responses take place behind the scenes as one of the steps in the establishment of a PPP connection. However, depending on how PAP is implemented in the client software, the user might still be required to provide an ID and password, which is then used in the negotiation process. The weaknesses of PAP are identical to

those of the simple password approach. However, if an administrator stores the password on the user's workstation the user does not have access to it and "social engineering" becomes much less effective. The downside to hiding passwords is that the administrator must keep track of the passwords for all machines with remote access capabilities. If the software on one of those machines is modified, it may be necessary to re-install the password.

CHAP

PAP transmits the password and user ID over the network as clear text, which could possibly be intercepted. Challenge Handshake Authentication Protocol (CHAP) eliminates this problem by using encryption to exchange messages during the PPP negotiation phase. Normally CHAP is implemented using an ID and password stored on the remote system, so it is not subject to the problems of the simple password approach, but it does require the same administrative effort as would a similar implementation of PAP.

Biometric

Although they have not yet gained wide acceptance, there are several authentication techniques that depend on the physical characteristics of the person requesting access. There are devices that can read fingerprints or look for unique patterns on a person's retina, and there are systems that can recognize a person according to voice characteristics.

SERVERS

To more easily manage a remote access facility, consider centralizing the authentication and logging process. In this approach, the terminal servers or access servers communicate with another system on the network that keeps track of authentication requirements for users and maintains appropriate logs.

Kerberos

Kerberos, an authentication protocol developed at the Massachusetts Institute of Technology (MIT), uses the Data Encryption Standard (DES) cryptographic algorithm for encryption and authentication. Kerberos, like other secret-key systems, is based on the concept of a trusted third party that performs secure verification of users and services. The Kerberos server is that trusted third party. Once the server has authenticated a user, it issues a temporary identification label referred to as a ticket. The user is then able to use the ticket rather than the user ID password combination to gain access to network resources. The lifetime of the ticket is configurable. The system administrator has the ability to issue short-lived tickets to less trusted users and tickets with longer lives to more trusted users.

In a remote access implementation, users provide user ID password combinations to the access server. This information is encrypted and sent to the Kerberos server, which returns a ticket to the access server. In the simplest case, the ticket would tell the terminal server to grant full access to the network, but it also could be used to control access to specific devices on the network. Any time the user requires access to a

device, the ticket is presented and the device checks with the Kerberos server to see if the ticket is valid. The Kerberos server allows or denies access accordingly.

Public domain versions of Kerberos are available from:

<<http://gost.isi.edu/info/kerberos/>>

TACACS+

Terminal Access Controller Access Control System plus (TACACS+) is a protocol that handles authentication, authorization, and accounting for remote access devices sold by Cisco Systems Inc. It uses TCP to transmit information between the remote access device and the computer that acts as a TACACS+ server. The information, with the exception of a short header, is encrypted. In addition to its capability of dealing with TCP/IP, TACACS+ also supports other remote access protocols such as:

- AppleTalk Remote Access (ARA)
- NetBIOS frame protocol control
- Novell Asynchronous Services Interface (NASI)
- Packet Assembler/Disassembler (PAD) connection

RADIUS

Remote Access Dial-In User Service (RADIUS) was developed by Lucent Technologies Inc. for use with their terminal servers. It has since been adopted by a number of other vendors, and it has become so popular that the RADIUS specification (RFC 2058) and RADIUS accounting standard (RFC 2059) are now proposed standard protocols. The text of the IETF proposed standards can be found at:

<<http://info.internet.isi.edu/in-notes/rfc/files/rfc2058.txt>>

<<http://info.internet.isi.edu/in-notes/rfc/files/rfc2059.txt>>

Like TACACS+, the RADIUS protocol is based on a client/server model. An access device passes user information to a designated RADIUS server and then acts on the response that is returned.

A RADIUS server can provide authentication and accounting services to one or more access devices. The server is responsible for receiving user connection requests, authenticating users, and then returning all configuration information necessary for the client to deliver service to the users. A RADIUS access server can be a dedicated workstation connected to the network or a separate program running together with other services in a multi-purpose computer.

RADIUS uses the UDP protocol as its communication mechanism, and it encrypts only password information. It can be configured to interact with other programs or servers, so it is possible to use RADIUS together with other protocols such as TACACS+ and Kerberos to provide centralized authentication authorization and accounting services for an entire organization.

Unrestricted, no-cost distributions of RADIUS are available from Lucent Technologies Inc. and Merit (University of Michigan) at the following URLs:

<<ftp://ftp.merit.edu/radius/releases/radius.3.5.6.basic.tar.{Z,gz}>>

<<ftp://ftp.livingston.com/pub/livingston/radius/>>

Other vendors have developed or adapted RADIUS server software to support their communications servers. The software is usually available directly from the vendors. However, in most cases, the RADIUS servers distributed by these companies are not considered supported products. They are provided as a convenience to the user.

There are now some commercial implementations of RADIUS with additional features not found in the standard distributions.

RAS

Users of Microsoft's Windows NT can attach one or more modems directly to a server or workstation and establish a remote connection with the server/workstation and therefore with other network resources. Authentication is accomplished with the same user ID and password that would be used from a computer that is directly connected to the network. Logging is accomplished through NT's event log.

Multi-port interface cards and a large bank of modems can be attached to a single NT server. This approach is normally not used, however, because it is more expensive and more difficult to manage than an access server that has been specifically designed for that purpose.

There are two major security weaknesses associated with attaching modems to individual servers or workstations. First, the machines with the modems are often managed by users who employ passwords that are easy to guess. A hacker dialling in may find this type of system is a much easier route into the network than a properly managed remote access server would be. The second problem could occur if a legitimate user dials into the network from another network that has an Internet connection. Once that dial-up connection is established, a hacker anywhere on the Internet has a direct connection through the dial-up user's network into the local network, effectively bypassing any firewalls or other controls that have been set up to restrict access from the Internet.

NetWare

A variety of third-party remote access products is available to work directly with a Novell server. Products include NetWare Loadable Modules (NLM), which work directly with Novell's NetWare Directory Services (NDS) to define user access and rights. There also are RADIUS implementations that link third-party hardware to NDS.

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

School jurisdictions have a particular concern about remote access security because their population includes people who consider circumvention of security measures a challenge. Such people can be particularly dangerous if they gain access to the network from a remote location where their actions can not be monitored or controlled. If authentication for remote access is obtained via simple passwords, implement a comprehensive password management strategy and take steps to prevent remote access to administrative resources from student-accessible facilities. Comprehensively log all remote access activities, and review the logs regularly.

REFERENCES

ISDN

<<http://www.isdnzone.com/dyndefault.htm>>

Modem Resources

<<http://www.modemshop.com/resources.html>>

Cable Modems

<<http://cabledatocomnews.com/cmic1.htm>>

Kerberos FAQ

<<http://www.faqs.org/faqs/kerberos-faq/general/>>

SECTION VI

CRACKERS AND HACKERS

Originally, clear distinctions were made between the terms "hacker" and "cracker." "Hacker" referred to people who were very familiar with computer hardware and software and used their expertise to develop sophisticated but non-threatening applications. Crackers were those who used their knowledge in a destructive way. "Hacker" is now commonly accepted as a description of someone who breaks into computer systems to accomplish unauthorized activities.

THE BASICS

Anyone can be a hacker. Tools and information are readily available, and opportunities are plentiful. School networks present particularly attractive targets for students because they do not have to worry about barriers protecting the network from external threats. Two techniques are particularly important for minimizing the effects of student hacking: administrative systems should be as well secured as possible, and student activities should be monitored as carefully as possible. Although the major threat is from students, school jurisdictions should not ignore the potential for hacking activity by staff.

THREATS

Hackers have varying levels of expertise and varying goals. For many, curiosity is the main reason for breaking into computer systems. Some want to see how far they can travel and what barriers they can cross. They are interested in improving their understanding of the inner workings of computers and networks, and in applying techniques that allow them to investigate the nuances of a particular type of computer system. Others are curious about information stored on the computer system. Students are interested in the records of fellow students. Employees might want to find out their colleagues' salaries. Even though the activities of the curious hacker may be relatively benign, they can still be responsible for major system disruptions or data loss if they make a mistake while exploring a computer system. If they come across information that should not be made public, they also have the potential to cause major problems for the organization.

For some hackers, profit is the motivation. They attempt to find information that may have some commercial value. Some operate like prospectors looking for gold. They invade potentially promising computer systems and try to find something of value. Others know exactly what they are looking for and direct their efforts toward one particular goal. One of the major tools of industrial espionage is invasion of computer systems. The people who indulge in this activity tend to be very competent, and they make every effort to carry out their exploits without being detected.

Some people view hacking as a game. They will invade a computer system and challenge the system administrators to chase them out and keep them out. They can

waste a great deal of administrative time, and they have the potential of causing system damage.

When "game-players" are in a losing position, they purposely destroy files or change system configurations as a sort of last-ditch gesture before abandoning the game and moving on.

The sole goal for some hackers is to cause system damage. They may have a grudge against the organization or a desire to cause problems for a business competitor. Or, they may simply enjoy destruction. These people are particularly dangerous, because their intention is to break into the system and cause as much trouble as possible as quickly as possible. Often, they accomplish their objective before they are even noticed, and the destruction is complete enough that tracing them becomes impossible.

HACKING TOOLS AND TECHNIQUES

The following web site is a typical example of a site that acts as a repository for hackers' tools and information:

<http://hem.passagen.se/argos/hacking.html>

This web site contains sections dealing with most of the popular operating systems, credit cards, the telephone system, viruses, and security tools. Anyone with a web browser can access this site and dozens of others like it to acquire tools that will help break into a computer system or network. Anyone with an inclination towards hacking can "start at the bottom" with these types of tools, improve their skills, and then develop contacts with more experienced hackers.

Less well-known bulletin board systems and mailing lists also exchange information useful to hackers. Hackers consider themselves members of a community. They pride themselves on their skills, and most of them enjoy sharing their accomplishments with their "colleagues."

The first step associated with breaking into a computer system is finding the computer. For computers attached directly to the Internet, this is usually an easy task. Tools such as the Internet Domain Name System, created for legitimate users, can be used by hackers. For computers not on the Internet, the most vulnerable points of attack are dial-up connections. Hackers can acquire software that will dial every telephone number within a specific range and identify where modems answer. Once they have a phone number or an Internet connection, hackers can attempt to connect by using common user ID password combinations found on many systems. They also may be able to fool the telephone system into letting them break into or take over a legitimate session.

After they have logged on to the computer, hackers will try to exploit any known loopholes in the system software that will allow them to gain more control. Often, they will install software that will capture user IDs and passwords from other system users. This software can work from within the operating system or may simply be an add-on to a program that is commonly used by users of the system. The objective is to collect enough information to allow them to obtain administrative access to the system. They also will employ software to see whether the administrator has discovered them.

Since hackers want to avoid being traced, a common tactic is to connect to a target system through a long string of systems that have already been compromised. They make great efforts to hide the fact that they are using those compromised systems as part of the path to the target. They try to cover their tracks by deleting information from log files and having as little impact as possible on the machines they are using as gateways.

Hackers interested in causing a disruption may launch a "denial-of-service" attack. These attacks make use of characteristics of the IP protocol or various operating system deficiencies to block access to a particular service or system. An example is the recent "ping of death." Someone discovered that very long ICMP echo requests directed over the network at certain computers could cause a failure in the operating system. The only way to recover was by restarting the system. This knowledge spread rapidly, and, as a result, many systems administrators had systems dying regularly for no particular reason. The manufacturers of the affected operating systems quickly developed fixes for the problem, but not before many machines had been attacked.

Some of the more sophisticated attacks involve manipulating the packets that carry information over the network. Tools are available that allow hackers to insert information into otherwise legitimate packets. This information can cause buffer overflows in the software of the target machine, and those overflows can open doors that allow the attacker to penetrate the system. Such attacks require much more capability than a simple password-guessing approach, but tools are available that make it possible for a relatively inexperienced hacker to mount a successful attack.

The Cuckoo's Egg by Clifford Stoll is a dated instructive account of a hacker's pursuits (The Bodley Head Ltd., © 1989, ISBN 0 370 31433 6).

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

Most organizations can acquire devices such as firewalls and remote access monitors to reduce the threat from external hackers. For a school jurisdiction, the biggest threat comes from within, and those external barriers, although important for other reasons, offer no protection. The best way to deal with these internal threats is to physically separate administrative and student activities, make sure that sensitive systems are configured to be as secure as possible, and be continually looking for suspicious activity. A good security policy can be very helpful in establishing guidelines for system security, and it also can provide a methodology for dealing with problems.

REFERENCES

Hacker Magazines

- <<http://www.2600.com/mindex.html>>
- <<http://www.phrack.com/>>

Hacker Sites

<<http://www.l0pht.com/>>

<<http://www.rootshell.com/>>

Hacking History

<<http://www.discovery.com/area/technology/hackers/hackers.html>>

The "Ping of Death" Page

<<http://www.dfm.dtu.dk/netware/pingod/ping.html>>

SECTION VII

FIREWALLS

THE BASICS

School jurisdictions should use firewalls to protect their networks from intrusion via the Internet. Firewalls log information on traffic from the local network to the Internet that provides statistical information and tracks possible misuse of the Internet connection. Firewalls are often used in conjunction with other applications such as filtering devices that prevent access to restricted web sites or newsgroups. Purchasing a firewall requires a considerable amount of research, and operating a firewall requires either an investment in training and time or an arrangement with an appropriate contractor. However, the relatively moderate effort and expense involved is a good investment. The potential costs of recovering data from unauthorized intrusion are overwhelming in comparison.

WHAT IS AN INTERNET FIREWALL?

The development of the Internet has been compared with the development of the American West in the nineteenth century. A few pioneers established a simple infrastructure that led to a massive influx of settlers and the development of a sophisticated society. During the early stages of that influx, there were few laws and even fewer law enforcement officers. Citizens had to depend on their own resources to protect themselves.

On the Internet today, unscrupulous people can roam about doing whatever they wish to facilities that have not been protected. Just as some citizens of the old West built fortifications to ensure their safety, the proprietors of computer networks connected to the Internet must establish mechanisms to secure their networks against outside intrusion. Such mechanisms are referred to as Internet firewalls. Firewalls can be simple and relatively easy to penetrate, or they can be complex, expensive, and very difficult to penetrate.

Although not everyone needs an impenetrable firewall, everyone needs some sort of protection. The degree of protection depends on the potential attraction to outsiders and the potential loss that could be suffered if the firewall were penetrated. Network administrators must balance their investment in protecting the network against the potential cost of damage arising from a successful attack.

FUNCTIONS PERFORMED BY FIREWALLS

Although the major purpose of an Internet firewall is to prevent damage from outside attacks, firewalls now perform a host of other duties as well. Simple firewalls accomplish some functions only in a limited way, whereas more complex firewalls handle a whole range of tasks with a great deal of sophistication. Some products combine all of the features described below with other popular Internet capabilities such as WWW, News, or DNS services. Generally, such hybrid devices are considered less

secure than special-purpose firewalls, but they may protect less vulnerable networks where cost saving is a major concern.

NETWORK PROTECTION

A firewall protects a network against damage from outside attacks by restricting outside connections to the devices on the protected network and by filtering any information that is transmitted from the outside to the inside.

CONTROLLED ACCESS

Even though it is desirable to isolate a vulnerable network from the Internet as fully as possible, there may be a need to allow specific individuals to gain access to the network. Firewalls accomplish this task by creating carefully controlled channels that allow only approved traffic. For example, the administrator of a network may wish to allow FTP transfers from a specific internal system to a group of identifiable systems on the Internet. With appropriate configuration, a firewall can allow only that specific activity while still protecting the network against other intrusions.

PRIVACY

Attackers who are unaware of exactly what computers are on a protected network or the topology of the network, find it much more difficult to plan and carry out an attack. By filtering out the types of packets that are used for network discovery and employing techniques such as IP address translation, a firewall can shield a protected network from prying eyes.

IP address translation can have an additional benefit; it allows administrators of the internal network to employ whatever address space they wish to use without worrying about how it will interact with the Internet. Many administrators with address-translating firewalls use non-routable IP addresses as described in RFC 1918 for their internal networks to take advantage of the benefits of having a wide range of addresses.

LOGGING

Since a firewall needs to inspect each packet that passes through, it also can record information about those packets. Such information can be used at a later date to trace attacks, detect misuse of the Internet by staff, or generate statistics.

STATISTICS

Log files or more sophisticated databases can be used to generate statistics concerning Internet usage. Such information as peak and average usage, most popular web sites, heaviest users, and busiest times of the day can be useful for planning facility upgrades or establishing usage policies.

POLICY ENFORCEMENT

Through logging, firewalls provide a passive mechanism for enforcing Internet usage policies. Most users who are aware that logging is taking place will not attempt to violate policies. They know that their activities will appear in the logs. If users do violate policies, the logs provide evidence of their activities. If more stringent controls are required, specific rules can be set up in many firewalls to prevent activities that are contrary to policy.

ISSUES AND PROBLEMS

Whenever a lock is put on a door, a series of new problems develops. Legitimate users must be provided with keys, the keys must be managed, it takes longer to go through the door because of the lock, etc. A similar set of problems arises as a result of a firewall installation. As with the lock, problems must be weighed against the added security that has been established.

RESTRICTION OF ACCESS FOR LEGITIMATE USERS

Firewalls have the potential of becoming irritating nuisances for users who have a legitimate need to access network resources. A good firewall minimizes the amount of interference a legitimate user faces while still acting as a secure barrier against unauthorized access.

POTENTIAL REDUCTION IN THROUGHPUT

Because firewalls must examine the traffic that passes through them, they affect the speed of connection to a network. The effect on speed can be offset with elaborate hardware, but that generally results in greater complexity and higher costs.

LITTLE PROTECTION FROM INSIDER ATTACKS

A lock on the door is not going to prevent someone who is already inside from causing problems. A firewall can not protect against unauthorized actions by someone who is directly connected to the protected network.

SINGLE POINT OF FAILURE

Most firewall implementations employ a single device to connect the protected network with the outside world. Although it is technically possible to create redundant firewall installations, they are often very complex and awkward to manage. In most practical situations, the firewall is the only route to the protected network, and a failure of this route will completely isolate the protected network.

REQUIRES MANAGEMENT

Even the simplest firewall requires a certain amount of management and a certain amount of expertise to design and implement appropriate rules. If logging is

implemented, additional effort is required to review the logs. If the user population or rule set is continually changing, even more management resources are required.

TYPES

PACKET FILTERING

The simplest devices that are marketed as firewalls do nothing more than examine the headers of IP packets and either pass or reject the packets according to a set of rules programmed into the firewall. Restrictions are generally based on source and destination IP addresses and port numbers. Such firewalls are adequate protection against unsophisticated attackers, but they are easily penetrated with readily available software that employs such techniques as IP spoofing or packet capture. Packet filters generally exhibit high throughput because the process for examining the packet headers is relatively simple.

ENHANCEMENTS TO PACKET FILTERING

One way to improve the effectiveness of packet filters is to consider each packet in the context of those that preceded it. Some firewall manufacturers have determined that the context of a stream of packets can help to determine whether the packets are legitimate. They have developed sophisticated algorithms that look not only at the individual packets, but also at their relationship to the other packets with which they are associated. If a packet appears which passes all of the rules associated with a single packet but does not fit into the expected context of the data stream, it is rejected. This approach results in a high data transfer rate, but it provides better protection than a simple packet filter.

PROXY SERVERS

Proxy servers are programs that reside on a device with two network interfaces. One interface is connected to the internal network, and the other is connected to the external network. The client program of a user on the internal network talks to the proxy server, and it in turn communicates with the server program on the external network (for example, a WWW server on the Internet). Depending on how they interact with the external systems, proxies may be considered to be application-level or circuit-level proxies. An application-level proxy interacts directly with the application protocol, understanding and interpreting the commands associated with that protocol. A circuit-level proxy simply creates a pathway between the client and server without understanding what information it is carrying. Application-level proxies are much more versatile and transparent to the user because they can take advantage of application-specific information to perform their tasks. However, they tend to be the least effective with respect to throughput because they interact extensively with each packet transmitted.

TOPOLOGY

In many cases, a single device is employed to protect a network, but some firewall implementations use two devices. This approach could be compared to building a

castle with a moat. The space between the two devices (analogous to the space between the moat and the castle wall) is referred to as a perimeter network or sometimes the DMZ (demilitarized zone). This network can be used as a location for servers that require a certain amount of protection but need to be accessed from the outside world (WWW or FTP servers, for example).

Another approach involves building a firewall with three interfaces. One interface is connected to the outside network, one to the protected network, and the third to a "service" network. The WWW and FTP servers are connected to the service network, and access to these devices is controlled by a set of rules on the firewall. The rules allow access to the service network from the outside, but they prevent penetration to the internal network.

RULES

AUTHENTICATION

Some firewalls require a user to provide a valid password before being allowed access to or from the protected network. The password could be a simple character string, or something more complex; for example, one of the single-use passwords generated by a proprietary authentication product such as SecureID. This system provides a fairly high level of control over user activities, but it places an extra impediment in the way of the user.

MAIL

One of the more common methods of attacking a network involves sending commands via otherwise innocuous-looking e-mail messages. Most of the current mail server software has built-in defenses against such attacks, but new attacks are continuing to appear. There also are a great number of servers still in operation with older, highly vulnerable mail software. Packet filtering firewalls can be programmed to reject e-mail packets, but such rejection may not be acceptable for an organization that wants protection but still needs to receive e-mail. The most common method for dealing with e-mail in a firewall involves installing a very simple mail server that does not respond to the types of command required to take control of the server. That server accepts mail from the outside world and then forwards it to the appropriate destination on the protected network.

DNS

The Internet Domain Name Service (DNS) is a useful tool for locating and communicating with networks and servers. Unfortunately, it also can be used by attackers to collect information about a target network that might help them penetrate the network. Most networks protected by a firewall tend to have a domain name server on the internal network to facilitate internal communication and a separate server on the external network to communicate with the world. The internal server is protected from external access by the firewall. The external server contains only information that is deemed to be acceptable for viewing by the outside world. Filters or a proxy allow the

internal server to query outside servers to resolve names for any outside locations that users may wish to access.

NEWS

News, transported by the NNTP protocol, has had significant problems with attacks. A typical firewall allows NNTP traffic from a specific source to a specific news server on the internal network. Further protection could be achieved by restricting the internal news server to that function alone. If it were successfully attacked, the damage would be restricted to news service (assuming that the attacker was unable to use the news server as a base from which to launch further attacks).

FTP

FTP servers have been successfully attacked on many occasions, so it is inadvisable to allow inbound access to an FTP server on the protected network. If FTP services to external users are needed, the FTP server should be placed outside the firewall, either directly on the Internet or on a separate perimeter network.

TELNET

The Telnet protocol was designed specifically to allow direct communication with the operating system of a remote computer. It provides an attacker with a means of directly interacting with a target machine. For this reason, it is highly inadvisable to allow unrestricted external access via Telnet to any computer on the protected network. Often operational requirements are that Telnet contact be made from the Internet to machines on the protected network. Restrict such contact as much as possible, enforce authentication with the firewall, and implement some sort of Telnet proxy on the firewall.

HTTP

Recently, attackers have been successfully exploiting weaknesses in HTTP servers. Often, the results of such exploits have been nothing more than embarrassing modifications to some of the web pages on the affected server. In some cases, though, a server has been used as a base for further attacks on the rest of the protected network. The precautions to take with respect to HTTP are similar to those described above for FTP.

OTHER PROTOCOLS AND SERVICES

Three protocols that are widely used for network management are ICMP (upon which PING is based), SNMP, and RMON. Because these protocols can allow an attacker to collect a tremendous amount of information about the organization of a network, they are generally blocked by firewalls. Normally, this does not cause much inconvenience for the users of a protected network, because they rarely need to use these protocols.

There is one exception. When an internal user attempts to use PING to check connectivity to an external host, the ICMP packets are blocked at the firewall, thereby causing the user to assume that the external host can not be contacted. This, of

course, prompts a call to the firewall manager with a complaint that the network is not working.

There is an ever-increasing demand to allow transmission of the many services available on the Internet through the firewall. Some services can be controlled and monitored with relative ease, while others present great opportunities for attackers. As firewall manufacturers try to keep pace with such developments, new proxies or filters are continually created. However, allowing any new service through a firewall should be done with considerable caution. If possible, some research should be carried out to ensure that the service does not represent a security risk, and to determine the best way to configure the firewall to provide optimal protection.

LOGS

A firewall log can be an extremely effective tool for security enhancement, but administrators also can be overwhelmed by the vast amount of information that may be collected.

INFORMATION RECORDED

Most logs collect information on the state of the firewall itself as well as information about traffic through the firewall. Some firewalls allow the manager to decide what information should be saved, while others do not. More extensive logs allow a much more thorough investigation of potential problems or break-in attempts, but they also present more of a management challenge in terms of review and storage.

Often, the process of logging information is tied directly to a process that raises alarms for critical incidents. These alarms may cause beeps on a console, send e-mail messages to an administrator, or even place a telephone call to a pager.

LOG MANAGEMENT

Review

The review of firewall logs can range from none at all to a detailed manual examination of every record. Most administrators have little time for this activity and depend on a variety of automated techniques, for example, the alarm filtering process described above. Another approach is to use a log analyzer program written specifically to scan a log and look for important information. Two such programs are TELEMADE <http://www.telemate.net> and logsurfer <http://www.cert.dfn.de/eng/team/wl/logsurf>. Often, administrators write special programs or filters to look for specific information in a log file; for example, the activity of a particular user in response to a complaint.

Archiving

There may be a need to review logs long after they have been captured. It is therefore advisable to establish an archiving process. For a busy site, archiving could involve storing a great deal of information so a process for automating this activity is usually established. Log files are regularly compressed and stored. The amount of time that

log files are kept depends in part on storage space. It is advisable to provide enough space to allow storage of at least seven days worth of logs and many organizations feel that logs should be kept for at least thirty days.

FIREWALL ACCESSORIES

Administrators often need control over information that comes into the network. Firewall manufacturers have attempted to fill this need by offering other services, either on the same hardware as is used for the firewall, or on associated systems. Such services include:

CONTENT FILTERS

Content filters prevent access to particular WWW sites or newsgroups or restrict the passage of information containing certain words, phrases, or other characteristics.

VIRUS CHECKERS

Virus checkers which detect viruses in incoming traffic (most often e-mail) should not be considered a substitute for a virus control policy. Viruses can infect a network from many sources in addition to external connections. Still, virus checkers are a very useful mechanism for limiting virus infection.

ENCRYPTION SERVICES

Many Internet users are using encryption techniques to prevent others from viewing their information. A somewhat more sophisticated implementation of encryption, the "virtual private network" (VPN), automatically encrypts information flowing among stations on an otherwise public network. Some firewall suppliers are adding software that can establish a VPN on the Internet among locations that have similar firewalls equipped with VPN functionality.

AUTHENTICATION SYSTEMS

Authentication systems focus more on who is using a facility than with the information being transmitted. Virtual private networks employ character strings that are associated with individuals. These strings, referred to as "keys," are used to encrypt and decrypt messages. Some firewalls that implement such features also have the capability of interacting with systems that manage such keys.

Some firewalls can be configured to force users to identify themselves before they transmit traffic through the firewall. These firewalls interact with proprietary authentication systems such as SecureID to accomplish this identification.

BUYING A FIREWALL

DEFINE REQUIREMENTS

Shopping for an automobile involves a wide array of possibilities and possible expenditures. Making the right decision can be difficult. Purchasing a firewall is even more daunting. A firewall shopper is like someone looking for a car who is being approached by people selling cars, skateboards, bicycles, freight trains, and jet planes—all claiming that they are selling automobiles. Many different devices with a great range of different capabilities are being marketed as firewalls. Before beginning, know exactly what you are looking for.

ESTIMATE REQUIRED THROUGHPUT

How many users on your network will be passing traffic through the firewall at the same time? What sort of traffic is it? What speed is the connection to the external network? What sort of response time will be required?

DECIDE UPON LEVEL OF SECURITY REQUIRED

Does the network need to be protected from all possible threats, or does it just need to be secure enough to deter casual hackers? Is user authentication required? What sort of information is needed for logs?

IDENTIFY OPERATIONAL REQUIREMENTS

Is ease of configuration important? What sort of alarm process is required? Is remote management necessary? What sort of log archiving process is required? How extensive must logging be? Are statistical reports needed? Are add-ons such as virus checkers or content filters needed?

SPECIFY PREFERRED OPERATING SYSTEM AND HARDWARE

Does your organization have a standard for hardware and operating systems?

Note: Although firewalls are dependent upon particular hardware and operating systems, they are generally designed as turn-key devices so that user interaction with the firewall is separated as much as possible from the underlying operating system. It is therefore much more important to choose a firewall on the basis of its characteristics and performance than upon its hardware and operating system. However, those criteria can be helpful in selecting among products with similar operating characteristics.

IDENTIFY POTENTIAL VENDORS

There are many companies who sell products called "firewalls", so an important task is to narrow down the field to a reasonable few. A little research through publications such as *Data Communications*, *Network World*, and *Network Computing*, can provide useful information for this process. Also, discussions with colleagues and Internet searches can yield some valuable information.

EVALUATE PERFORMANCE AND EASE OF USE

Several publications regularly review firewalls and publish the results. Such reviews can provide a great deal of useful information, but the ratings should not necessarily be considered valid for all circumstances. Because firewalls perform many tasks, and because they come in many shapes and sizes, most reviews are not extensive enough to provide a complete survey. In fact, the results of some reviews may be misleading because they concentrate on a few specific characteristics such as throughput or the capability of blocking a few particular types of attack. However, such reviews, when combined with input from other sources such as colleagues and vendors can be very useful.

To evaluate the ease of managing a firewall, be sure to request a hands-on session. Many manufacturers provide evaluation copies of their products or have demonstration systems at their office. In addition to taking a hands-on "test drive," consult with other users of the systems being evaluated.

EVALUATE SUPPORT AND UPGRADEABILITY

If internal resources are available to support the firewall, ask the vendor for information about available training. If such resources are not available, the issue of local support becomes important. Can the vendor or a contractor supply such support? How much will it cost? How about telephone support?

New attack methods are constantly appearing and new services continue to be developed. Therefore, the vendor's policy on the cost and availability of software revisions is important. How often are new revisions issued? How easy are they to install? Does the vendor provide patches to deal with specific problems between revisions? If so, how easy are they to acquire and install? What is the cost of ongoing software support?

FIREWALL ADMINISTRATION

Once it is set up, a simple packet filtering firewall does not require a great deal of administration unless rules need to be changed. However, the security provided by such a device is limited. The administration required for more sophisticated firewalls is increased because of the number of services and features available. However, in general, the actual administration process is simpler because the manufacturers of such devices normally include more sophisticated tools.

A firewall administrator can expect to undertake daily tasks such as checking the logs, responding to alarms, and monitoring the log archiving process. Other activities include monitoring disk space, periodic software updates, rule modification in response to user requirements, and managing user lists for firewalls that employ authentication.

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

School jurisdictions have some different requirements regarding firewalls. To accommodate very high usage during peak periods school jurisdictions need a firewall with relatively high throughput. Since school jurisdictions also have a certain number of users who view security measures as a challenge, the firewall has to be robust and immune to attacks from the internal network. School jurisdictions tend not to have the resources to devote to appropriate staff for firewall implementation and administration and must therefore choose a firewall that can be managed by existing staff or have access to appropriate outside resources to manage the firewall. Most school jurisdictions need a firewall that supports content filtering.

REFERENCES

Firewall Resource List

[<http://www.cs.purdue.edu/coast/firewalls>](http://www.cs.purdue.edu/coast/firewalls)

Internet Firewalls Frequently Asked Questions

[<http://www.clark.net/pub/mjr/pubs/fwfaq/>](http://www.clark.net/pub/mjr/pubs/fwfaq)

Searchable Firewalls Mailing List Archive

[<http://www.nexial.nl/cgi-bin/firewalls>](http://www.nexial.nl/cgi-bin/firewalls)

Commercial Firewall Vendors

[<http://www.waterw.com/~manowar/vendor.html>](http://www.waterw.com/~manowar/vendor.html)

Firewalls Mailing List

The Firewalls mailing list is for discussions of Internet firewall security systems and related issues. Relevant topics include the design, construction, operation, maintenance, and philosophy of Internet firewall security systems. The Firewalls mailing list is hosted by GNAC, Inc.

The list is fairly high volume (often 20+ messages per day, sometimes 100+), and is available in both regular and digest form. The digest has all the same messages as the regular list, just collected into digest form (so that you get a few lengthy messages each day, rather than lots of short ones).

To subscribe to the regular mailing list, send the following command in the body of an email message (NOT on the "Subject:" line) to majordomo@lists.gnac.net
subscribe firewalls

To subscribe to the digest, send the following command in the body of an email message (NOT on the "Subject:" line) to majordomo@lists.gnac.net
subscribe firewalls-digest

SECTION VIII

APPLICATIONS

THE BASICS

Every application employed on a computer network has its own unique security risks and requirements. All of these risks and requirements must be evaluated before the system is purchased. Each application has its own unique focus and must therefore be researched separately. However, three applications—electronic mail, the World Wide Web, and single sign on—are sufficiently generic to be discussed in this report.

Electronic mail is one of the most commonly used applications on the majority of computer networks. It also is one of the most commonly abused applications. A secure e-mail system should not allow users to send forged messages, and it should not allow the transmission of unwanted messages to large numbers of users. Also, there should be a means of adding a digital signature to assure the recipient of the source of the message.

Security for world-wide web applications involves two central issues: the restriction of access to particular sites, and the secure transmission of information.

Access restriction can be accomplished with a device called a proxy server that sits between the local area network and the Internet. These servers intercept requests for connections to web sites, and verify the requests against a list of forbidden sites. Only connections to permitted sites are established.

Methodologies for secure transmission of information involve the use of encryption, digital signatures, and trusted third-party authorities who verify the authenticity of a particular user or location.

One way of simplifying password management when a network incorporates a variety of different computer systems is to install a separate "single sign-on" product that is responsible for controlling access to all systems on the network. Such systems enhance security by incorporating encryption and minimizing the chance that a user will leave a written copy of a password in an obvious place.

E-MAIL

PROTOCOLS

SMTP

Simple Mail Transfer Protocol (SMTP) is the foundation of Internet mail. A server wishing to send mail via SMTP communicates with the recipient's server using a series of commands and responses. Because there is no guarantee that a connection will always be established between the sending and receiving server, the software that employs SMTP normally holds on to a message until a connection can be established. There also is a capability for automatically forwarding mail from one server to another,

so that users can have their mail automatically redirected, or one server can act as a backup for another. In that situation, the backup server receives and stores mail until the primary server becomes available and then forwards the mail to the primary server.

POP

Post Office Protocol (POP) allows mail to be efficiently and quickly transferred from a central mail server to a client workstation. POP is most often used when a user wishes to minimize connect time to the server and view received mail on the workstation. Most popular workstation-based e-mail programs employ POP to download mail from the server. Normally, the mail is deleted from the server once it has been downloaded. There is an option that allows a copy of the downloaded mail to remain on the server, but it is not foolproof. Problems are common with synchronizing the mail file on the server and the mail on the client workstation.

POP3 supports two possible means for a client to identify itself to a server: a user ID/password combination or a somewhat more secure method called Authenticated Post Office Protocol (APOP) that makes use of a shared key to encrypt password information that passes from the client to the server.

IMAP

The Internet Message Access Protocol (IMAP) has been described as a functional superset of POP. It can be used in the same way as POP for transferring messages from a server to a client, but it also has other powerful features. The major difference between the protocols is due to the fact that IMAP allows messages to be stored on the server in user-specified files. This is a particularly important feature for users who are mobile. They can view their mail from any location and are not required to download mail to the workstation they are using. IMAP allows concurrent access to a shared mailbox from multiple platforms, concurrent access to mailboxes on multiple servers, and the selective transfer of messages from the server to any client. For messages that contain multiple MIME-encoded attachments, IMAP allows selective fetching of those attachments.

Two versions of IMAP are currently in common use. The protocols are described in detail at:

<<http://www.faqs.org/rfcs/rfc1176.html>> and
<<http://www.faqs.org/rfcs/rfc2060.html>>

Most implementations of IMAP use a clear text user ID/password authentication mechanism. A recently developed protocol called Challenge-Response Authentication Mechanism (CRAM) functions in a similar way to APOP, but it has an additional advantage. CRAM does not require the client and the server to have the shared key for password encryption stored as clear text.

IMSP

Internet Messaging Support Protocol (IMSP) was developed as an extension to IMAP to allow clients to automatically deal with multiple servers. IMSP provides a central repository for information on address books, user preferences, and other options that can be accessed from any client. Work on IMSP has now largely been redirected toward another standard called Application Configuration Access Protocol (ACAP) that is a more feature-rich implementation intended for general-purpose use.

PROPRIETARY PROTOCOLS AND SERVICES

Many network-based e-mail systems employ proprietary protocols and services, each system has its own strengths and weaknesses, and its own security problems and solutions. If the system communicates via the Internet, it almost certainly has an SMTP gateway that interacts with the rest of the world using the SMTP protocol.

E-MAIL SECURITY

It is easy to send an e-mail message with a false sender name and address. So one of the issues to consider with respect to e-mail security is a means of verifying that the message actually came from where it appeared to originate. The most effective method involves digital signatures. The sender of the message attaches an encoded character string that is decoded by the recipient. The string is based on the contents of the message and is encoded using a "private key" known only to the sender. The recipient uses a "public key," which matches the sender's private key to decode the signature and then compares the signature with a character string generated from the message by the same means as the signature was generated by the sender. If the two strings match, the recipient knows that the sender was genuine, and the message was not altered during the transmission.

The same technology, referred to as "public key encryption," can be used to encrypt the entire contents of the message and prevent its contents from being viewed during transmission.

One other piece of information also can be attached to a message to verify its authenticity. A "digital time stamp" is an encrypted representation of the time the message was sent.

In the past, the SMTP protocol has been used as a means of attacking UNIX computer systems. Earlier implementations of mail software offered many features for communicating with the server that were convenient but also open to abuse. Means have been found to prevent such abuse in most situations, but as new versions of e-mail software are developed to enhance protection, the hacker community continues to develop new methods of compromising the software. Nevertheless, the most important e-mail security procedure is to ensure that the latest version of the software is installed and new security patches implemented as soon as possible.

Because of the design of SMTP and the general philosophy about storing and forwarding e-mail upon which Internet mail is based, it is relatively easy to use an unsuspecting server as a broadcast host for unsolicited mail. Developers of e-mail

server software are now incorporating features to prevent this type of abuse, but even some of the latest releases do not offer a foolproof solution.

Because of the store-and-forwarding nature of e-mail, it is possible to completely shut down a mail server by sending it many large messages that it can not forward immediately. Eventually, the available storage space fills up, and no more mail can be received. The only effective means of dealing with this problem is to regularly monitor available space, and possibly place restrictions on the amount of space allocated to any individual user.

SPAMMING

The practice of sending unsolicited mail to a large number of recipients, referred to as "spamming," can have negative effects on the user who receives the unwanted message, as most Internet mail users have discovered. It can have even worse effects on administrators of the server from which the spam originated as they try to deal with complaints from thousands of unhappy recipients.

Some of the newer e-mail software has built-in features called "spam filters" that attempt to protect against unsolicited mail from outside the network. These filters can be configured to reject mail from particular IP addresses or particular users, but since it is possible to fake a sender's address, and use an intermediate (and often unsuspecting) server as a mail forwarder, such filters have only limited capabilities.

Often, mail software can be tailored to prevent outgoing spam messages by limiting the number of addresses to which a message may be sent, and by logging information about the user attempting to send the messages.

CLIENT SOFTWARE

Most e-mail client software for non-proprietary implementations uses SMTP as the mechanism for sending messages, and either POP or IMAP as the mechanism for receiving messages from a server.

The security of e-mail messages on a client workstation depends on how secure the workstation itself is. Appropriate workstation security measures must be in effect to complement whatever e-mail security measures are in place.

DIRECTORY SERVICES

X.500

One of the challenges associated with implementing a large computer network is keeping track of all of the users of that network. X.500 is a standard developed for creating extremely comprehensive directories that intercommunicate with one another to share information. A valuable feature for an e-mail system is the capability of interacting with X.500 directories to locate and manage e-mail addresses.

LDAP

The Lightweight Directory Access Protocol (LDAP) is a simplified mechanism for working with information in X.500 directories. Because LDAP uses a simplified data representation for the various protocol elements, clients are smaller, faster, and easier to implement than a full-scale X.500 client.

Currently, there are no official standards for applying security constraints to LDAP, although various manufacturers have added extensions to enhance security. The current thinking about access control and LDAP is presented at

<<http://www.ietf.org/internet-drafts/draft-ietf-ldapext-acl-reqts-00.txt>>

WWW

APPLETS

"Applets" are one of the more controversial issues in the area of WWW security. These computer programs are downloaded from a web site to a user's computer and then executed on the user's computer. Applets differ from "plug-ins" in one major respect. A plug-in such as *Shockwave* or *RealAudio* must be specifically downloaded and installed. Applets, on the other hand, are downloaded automatically when a browser with the appropriate features accesses a web site that employs applets.

There are two major technologies for creating and executing applets: Java, developed by Sun Microsystems, and Active-X from Microsoft. From the user's point of view, the biggest difference between the two revolves around security. Java has supposedly been designed to prevent applets from interfering with or accessing private data on the user's computer. According to its developers, security is an automatic feature that should not require any interaction with the user. Microsoft, on the other hand, counts on the user to decide whether or not an applet can be trusted. Before executing the applet, the user is asked to approve its use, based on whether the user feels that the site from which the applet originated is trustworthy.

Problems have been detected with both technologies, so some manufacturers of web-filtering products have added features that place restrictions on incoming applets. To date, the problems have been well documented, and in most cases have been of the nuisance variety rather than significant security breaches. Because applets can add so much to the functionality of a web site, an across-the-board restriction will probably result in user discontent. The best approach is to maintain awareness of potential problems and deal with them as required.

SERVERS

Operators of web servers are at much greater risk of security problems than those who use browsers. Servers can be accessed by anyone on the network, and because server software tends to be much more complex than browsers, it tends to have more bugs and therefore more potential security holes.

If a server is simply providing information to members of the local network, the best and simplest way to protect it is to locate it on the internal network behind a firewall. Another server, providing strictly public information, can be located outside the firewall.

If the external server is compromised, the only loss will be the public information, which can be replaced easily. However, if the internal server is compromised by someone on the internal network, it could provide easy access to the rest of the network. The firewall alone does not provide complete protection for the internal server or the network. Appropriate measures to minimize the risk to the internal server include all of the general steps required for protecting any server as well as additional measures appropriate to the WWW server software being employed. An up-to-date list of potential security problems and appropriate solutions can be found at

<http://www.w3.org/Security/Faq/www-security-faq.html>

Another area of concern is the use of scripts executed on the server by the WWW server software. Many such scripts are available on the Internet to provide added features to a web application. Before implementing any such script, however, have someone who is familiar with the scripting language and the security risks of the language review it for potential security problems.

SECURE TRANSMISSION OF INFORMATION

Various encryption schemes are available for providing secure transmission of information via the World Wide Web. The two most popular are Secure Sockets Layer (SSL), developed by Netscape Communications Corporation and Secure HTTP (S-HTTP), which is being supported by CommerceNet, a coalition of businesses whose goal is to develop commercial uses for the Internet.

SSL includes provisions for authentication of both client and server as well as a mechanism for encrypting data in transit. It is used quite widely, and has been implemented on most popular browsers and servers. SSL works at a low level in the communications software, so it can be employed with various high level protocols such as FTP and NNTP as well as HTTP.

S-HTTP is a higher-level protocol than SSL. It works only with HTTP, but it has greater potential for extensions than SSL.

The sender and recipient of an encrypted message must employ unique strings of digits called keys. Generally, an individual has both a private key and a public key. The private key is used to encrypt the message, and the public key is used to decrypt it. Another use of such keys is to allow a user who is interacting with a server to be sure that the server is actually at the site it purports to serve. To accomplish this, various organizations have established themselves as "Certificate Authorities." For a fee, these organizations provide a private and a public key and vouch for the accuracy of the keys to anyone who wishes to ask.

SINGLE SIGN ON

Simplified user authentication is one of the advantages of standardizing to a particular type of computer system. Generally, once users log on to the network they can access any authorized resource on the network without further authentication. Members of large organizations with many different computer systems often find they must maintain several different user names and passwords so they can log on to each system, and administrators must manage a separate user list and authentication mechanism for each system.

This approval can be a major problem for the users, particularly if password-aging standards require the user to change passwords regularly. Because it is extremely difficult to remember a large number of regularly changing passwords, users tend to use the same password on multiple systems or keep a written record of their passwords. Both of these approaches represent security risks. As well, administrators are continually having to reset passwords for users who have forgotten or lost them. To make life easier for both users and administrators, several manufacturers are now developing products that allow users to employ a single log-on sequence to connect to multiple systems.

Single sign-on systems employ a special client program at each workstation. Users who wish to connect to the network initiate the client program. That program requests a user name and password, and then carries on a dialogue with a central registry on the network that authenticates the user. The central registry returns a list of systems for which the user is authorized and a process for connecting to each of those systems. Generally, all of this information is encrypted so that it can not be intercepted and misused. When users wish to interact with a particular system on the network, they select that system from a menu displayed by the client, and the client carries on an authentication dialogue with the selected system on behalf of the user.

UNIQUE REQUIREMENTS OF SCHOOL JURISDICTIONS

School jurisdictions often use proprietary e-mail systems that incorporate features that are of particular use in the educational environment. The security issues outlined in the above discussion may or may not apply to such systems. It is important to review the mail system in use, carry out research on potential security problems, and discuss problems with the manufacturer of the system.

Many school jurisdictions consider web site blocking essential to prevent students from reaching "unsavoury" sites. Several things should be considered when evaluating a product for such purposes. The proxy server and the associated software must be sufficiently powerful to prevent the site blocking process from becoming a restriction on throughput capability to the Internet. Because web sites are constantly changing, there must be a means to regularly update the list of forbidden sites. There also should be an easy way to add to or subtract from the list, and there should be a logging process so that activity can be reviewed easily if necessary.

Public key encryption technology can be useful for school jurisdictions that exchange information with external users via the Internet. It also can be a valuable mechanism for

authenticating internal users who may be accessing internal web-based applications. With this approach, the school jurisdiction could set up its own certificate service and assign keys to each of its users for use only on the internal network.

The need for a "single sign-on" product will depend on how many different computer systems are in use on the network and how many users require access to more than one system.

REFERENCES

SMTP Standard

[<http://www.faqs.org/rfcs/rfc821.html>](http://www.faqs.org/rfcs/rfc821.html)

POP Standard

[<http://www.faqs.org/rfcs/rfc1939.html>](http://www.faqs.org/rfcs/rfc1939.html)

IMAP Standards

[<http://www.faqs.org/rfcs/rfc1176.html>](http://www.faqs.org/rfcs/rfc1176.html)

[<http://www.faqs.org/rfcs/rfc2060.html>](http://www.faqs.org/rfcs/rfc2060.html)

CRAM

[<http://www.faqs.org/rfcs/rfc2095.html>](http://www.faqs.org/rfcs/rfc2095.html)

ACAP

[<http://www.faqs.org/rfcs/rfc2244.html>](http://www.faqs.org/rfcs/rfc2244.html)

LDAP

[<http://www.faqs.org/rfcs/rfc1777.html>](http://www.faqs.org/rfcs/rfc1777.html)

Security Analysis of Java

[<http://www.cs.princeton.edu/sip/java-faq.html>](http://www.cs.princeton.edu/sip/java-faq.html)

Java vs. ActiveX

[<http://www.cs.princeton.edu/sip/java-vs-activex.html>](http://www.cs.princeton.edu/sip/java-vs-activex.html)

WWW Security FAQ

[<http://www.w3.org/Security/Faq/www-security-faq.html>](http://www.w3.org/Security/Faq/www-security-faq.html)

SSL

[<http://home.netscape.com/newsref/std/SSL.html>](http://home.netscape.com/newsref/std/SSL.html)

S-HTTP

[<http://www.terisa.com/shttp/current.txt>](http://www.terisa.com/shttp/current.txt)

Public Key Cryptography

<<http://www.cis.ohio-state.edu/text/faq/usenet/cryptography-faq/part06/faq.html>>

SECTION IX

CLOSING COMMENT

In educational environments, where an exceptionally large number of users are accessing different personal computers at different times and in different locations, the challenge of controlling and monitoring access is huge. Because of this fact, schools potentially face greater security risks than private businesses.

It is absolutely essential, therefore, for school jurisdictions that are using wide area networks with connections beyond the organization to research, plan and regularly update their security policies and procedures. The potential results of inadequate security include:

- inappropriate use of confidential or private information,
- having databases and/or files modified, damaged or destroyed, and
- interruptions in service, with tremendous amounts of learning, teaching and administrative time being lost.

The key considerations in developing a network security system are technical (what is available and practical to use), budgetary (maintenance and operation as well as capital costs) and administrative (who will keep the security system up and running and how).

This document is designed to introduce school jurisdiction staff to basic information about the physical aspects of network security (products, technical solutions, appropriate systems) and the human aspects (audits and evaluation, communication, staff training, and so on).

GLOSSARY OF TERMS

The source for this glossary is the site PCWebopedia <<http://www.pcwebopedia.com>>. Some of the entries have been adapted for the purpose of this document.

ACAP	Short for <i>Application Configuration Access Protocol</i> , an electronic mail protocol being developed by the IETF (Internet Engineering Task Force) to complement IMAP4 (the latest version of IMAP). ACAP supports related e-mail services such as subscribing to bulletin boards, and organizing and searching mailboxes and address books.
ActiveX	A loosely defined set of technologies developed by Microsoft. When talking about ActiveX, most people commonly think only of ActiveX controls. ActiveX controls facilitate the implementation and use of Active X technologies on the Internet.
APOP	Short for <i>Authenticated Post Office Protocol</i> . A version of POP which allows for authentication of the user.
Applets	A program designed to be executed from within another application. Unlike an application, applets can not be executed directly from the operating system.
ARP	<i>Address Resolution Protocol</i> converts Internet Protocol addresses, a computer's address on the Internet, into physical addresses, a computer's unique identifier on a network.
Backbone	Another term for bus, the main wire that connects nodes (computers or devices on a network). The term is often used to describe the main network connections composing the Internet.
Bandwidth	The amount of data that can be transmitted in a fixed amount of time. For digital devices, the bandwidth is usually expressed in bits per second (bps) or bytes per second. For analog devices, the bandwidth is expressed in cycles per second, or Hertz (Hz).
BDC	Short for <i>Backup Domain Controller</i> . Contains the same information as the PDC and is primarily used to share the load with the PDC in authenticating all logon requests.
Biometric	The process of using physical characteristics of a user to authenticate users. By using devices connected to the computer, features such as fingerprints or retina identification can be used. This is not yet as widely accepted as PAP or CHAP.
BIOS	Pronounced "bye-ose," an acronym for <i>basic input/output system</i> . The BIOS is built-in software that determines what a computer can do without accessing programs from a disk. On PCs, the BIOS contains all the code required to control the keyboard, display screen, disk drives, serial communications, and a number of miscellaneous functions.

Bridge	A device that connects two local area networks (LANs), or two segments of the same local area network (LAN). The two LANs being connected can be alike or dissimilar. For example, a bridge can connect an Ethernet network to a Token-Ring network.
Bus	In networking, a bus is a central cable that connects all devices on a local area network. It also is called the backbone.
Cable Modem	A modem designed to operate over cable TV lines. Because the coaxial cable used by cable TV provides much greater bandwidth than telephone lines, a cable modem can be used to achieve extremely fast access to the World Wide Web. This, combined with the fact that millions of homes are already wired for cable TV, has made the cable modem a very strategic and valuable asset for Internet and cable TV companies.
Cache	Pronounced <i>cash</i> , a special high-speed storage mechanism. It can be either a reserved section of main memory or an independent high-speed storage device.
CERT	Short for <i>Computer Emergency Response Team</i> . The CERT Co-ordination Center (CERT/CC) is located at the Software Engineering Institute (SEI), a federally funded research and development center at Carnegie Mellon University in Pittsburgh, Pennsylvania. Following the Internet Worm incident, which brought ten per cent of Internet systems to a halt in November 1988, the Defense Advanced Research Projects Agency charged the Software Engineer Institute with setting up a centre to co-ordinate communication among experts during security emergencies and to help prevent future incidents. Since then, the CERT/CC has helped to establish other response teams while maintaining leadership in analyzing vulnerabilities and threats. The CERT web site can be found at < http://www.cert.org >.
CHAP	Short for <i>Challenge Handshake Authentication Protocol</i> , a type of authentication in which the authentication agent (typically a network server) sends the client program a key to be used to encrypt the username and password. This enables the username and password to be transmitted in an encrypted form to protect them against eavesdroppers. Contrast with PAP.
CIAC	Short for <i>Computer Incident Advisory Capability</i> . The CIAC provides on-call technical assistance and information to Department of Energy (DOE) sites faced with computer security incidents. This central incident handling capability is one component of all encompassing service provided to the DOE community. The other services CIAC provides are: awareness, training, and education; trend, threat, vulnerability, data collection and analysis; and technology watch. This comprehensive service is made possible by a motivated staff with outstanding technical skills and a customer service orientation. CIAC is an element of the Computer Security Technology Center which supports the Lawrence Livermore National Laboratory.
CLI	Short for <i>Calling Line Identification</i> . The same technology that gives telephone users call display also can be used to screen phone calls into remote access servers from unauthorized locations.

CPU	Short for Central Processing Unit . The CPU is simply the brains of the computer.
Data Transfer Rates	The speed with which data can be transmitted from one device to another. Data rates are often measured in megabits (million bits) or megabytes (million bytes) per second. These are usually abbreviated as Mbps and MBps, respectively.
DHCP	Short for Dynamic Host Configuration Protocol , a protocol for assigning dynamic IP addresses to devices on a network. With dynamic addressing, a device can have a different Internet Protocol (IP) address every time it connects to the network. In some systems, the device's IP address can even change while it is still connected. DHCP also supports a mix of static and dynamic IP addresses.
Direct Access	A modem installed on a computer on the network allows users to log in remotely by dialling into that modem using a modem and regular phone lines. Depending upon capabilities, the user can then have access to resources on the network.
DNS	Short for Domain Name System (or Service) , an Internet service that translates domain names into Internet Protocol (IP) addresses. Because domain names are alphabetic, they are easier to remember. The Internet however, is really based on IP addresses. Every time you use a domain name, therefore, a DNS service must translate the name into the corresponding IP address. For example, the domain name www.example.com might translate to 198.105.232.4.
DSP	Short for Digital Signal Processing , which refers to manipulating analog information, such as sound or photographs that have been converted into a digital form.
DSP-Based Access Servers	The latest servers for remote access make use of a technology called digital signal processing. With these devices, the functions of the modems and the terminal server are combined into a single device.
ERA Model	Short for Environment, Risk, and Assurance Model in Novell network systems. A model used by Novell to help design and implement a workable security policy.
Ethernet	A local area network (LAN) protocol developed by Xerox Corporation in co-operation with DEC and Intel in 1976. It is one of the most widely implemented LAN standards.
FAT	Short for File Allocation Table . A table that the operating system uses to locate files on a disk. Due to fragmentation, a file may be divided into many sections that are scattered around the disk. The FAT keeps track of all these pieces.
Finger	A UNIX program that takes an e-mail address as input and returns information about the user who owns that e-mail address. On some systems, finger only reports whether the user is currently logged on. Other systems return additional information, such as the user's full name, address and telephone number.

Firewall	A system designed to prevent unauthorized access to or from a private network. Firewalls can be implemented in both hardware and software, or a combination of both. Firewalls are frequently used to prevent unauthorized Internet users from accessing private networks connected to the Internet, especially Intranets. All messages entering or leaving the Intranet pass through the firewall, which examines each message and blocks those that do not meet the specified security criteria.
FTP	Short for <i>File Transfer Protocol</i> , the protocol used on the Internet for sending files.
HTTP	Short for <i>HyperText Transfer Protocol</i> , the underlying protocol used by the World Wide Web. HTTP defines how messages are formatted and transmitted, and what actions Web servers and browsers should take in response to various commands.
Hubs	A common connection point for devices in a network. Hubs are commonly used to connect segments of a local area network (LAN). A hub contains multiple ports. When a packet arrives at one port, it is copied to the other ports so that all segments of the LAN can see all packets.
ICMP	The <i>Internet Control Message Protocol</i> was designed in order to report a failure in routing back to the source of any packet being routed.
IMAP	Short for <i>Internet Message Access Protocol</i> , a protocol for retrieving e-mail messages..
IMSP	Short for <i>Internet Messaging Protocol</i> . An extension to IMAP developed to allow clients to deal with multiple servers. Work on IMSP has been largely redirected to a new standard called ACAP.
IP	Abbreviation of <i>Internet Protocol</i> , pronounced as two separate letters. IP is something like the postal system. It allows you to address a package and drop it in the system, but there is no direct link between you and the recipient. See also TCP/IP.
IPX	Short for <i>Internetwork Packet Exchange</i> , a networking protocol used by the Novell NetWare operating systems.
ISDN	Abbreviation of <i>Integrated Services Digital Network</i> , an international communications standard for sending voice, video, and data over digital telephone lines. ISDN requires special metal wires and supports data transfer rates of 64 Kbps (64,000 bits per second). Most ISDN lines offered by telephone companies give you two lines at once, called B channels. You can use one line for voice and the other for data, or you can use both lines for data to give you data rates of 128 Kbps, three times the data rate provided by today's fastest modems.
ISP	Short for <i>Internet Service Provider</i> , a company that provides access to the Internet. For a monthly fee, the service provider gives you a software package, username, password and access phone number. Equipped with a modem, you can then log on to the Internet and browse the World Wide Web and send and receive e-mail.

Java	A general purpose programming language with a number of features that make it suitable for use on the World Wide Web.
Kbps	Short for <i>Kilobits per second</i> , a measure of data transfer speed. Modems, for example, are measured in Kbps. Note that one Kbps is 1000 bits per second, whereas a KB (kilobyte) is 1024 bytes. Data transfer rates are measured using the decimal meaning of K whereas data storage is measured using the powers-of-2 meaning of K. Technically, kbps should be spelled with a lowercase k to indicate that it is decimal but almost everyone spells it with a capital K.
Kerberos	An authentication system developed at the Massachusetts Institute of Technology (MIT). Kerberos is designed to enable two parties to exchange private information across an otherwise open network. It works by assigning a unique key, called a <i>ticket</i> , to each user that logs on to the network. The ticket is then embedded in messages to identify the sender of the message.
Key	A password or table needed to decipher encoded data.
LAN	Short for <i>Local Area Network</i> . A computer network that spans a relatively small area. Most LANs are confined to a single building or group of buildings. Most LANs connect workstations and personal computers. Each node (individual computer) in a LAN has its own central processing unit with which it executes programs, but it also is able to access data and devices anywhere on the LAN. This means that many users can share expensive devices, such as laser printers, as well as data. Users also can use the LAN to communicate with each other, by sending e-mail or engaging in chat sessions.
LDAP	Short for <i>Lightweight Directory Access Protocol</i> , a set of protocols for accessing information directories. LDAP is based on the standards contained within the X.500 standard, but is significantly simpler. And unlike X.500, LDAP supports TCP/IP, which is necessary for any type of Internet access. Because it is a simpler version of X.500, LDAP is sometimes called X.500-lite.
MAC address	Short for <i>Media Access Control address</i> , a hardware address that uniquely identifies each node of a network.
Mbps/MBps	see <i>Data Transfer Rates</i> .
NDS	Short for <i>Novell Directory Services</i> , the directory services for Novell Netware networks. NDS complies with the X.500 standard and provides a logical tree-structure view of all resources on the network so users can access them without knowing where they are physically located. NDS also interoperates with other types of networks.
NIC	Short for <i>Network Interface Card</i> . Often abbreviated as <i>N/C</i> , an expansion board you insert into a computer so the computer can be connected to a network. Most NICs are designed for a particular type of network, protocol, and media, although some can serve multiple networks.
NNTP	Short for <i>Network News Transport Protocol</i> , the protocol used to post, distribute, and retrieve USENET messages. The official specification is RFC 977.

NOS	Short for Network Operating System . An operating system that includes special functions for connecting computers and devices into a local area network.	
NTFS	Short for New Technology File System , one of the file systems for the Windows NT operating system (Windows NT also supports the File Allocation Table file system). NTFS has features to improve reliability, such as transaction logs to help recover from disk failures. To control access to files, you can set permissions for directories and/or individual files. NTFS files are not accessible from other operating systems such as DOS.	
OSI	Short for Open System Interconnection , an ISO (International Organization for Standardization) standard for worldwide communications that defines a networking framework for implementing protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station, proceeding to the bottom layer, over the channel to the next station and back up the hierarchy.	
OSI REFERENCE MODEL		
Layer	Name	Function
7	Application Layer	Program-to-program communication.
6	Presentation Layer	Manages data representation conversions. For example, the Presentation Layer would be responsible for converting from EBCDIC to ASCII.
5	Session Layer	Responsible for establishing and maintaining communications channels. In practice, this layer is often combined with the Transport Layer.
4	Transport Layer	Responsible for end-to-end integrity of data transmission.
3	Network Layer	Routes data from one node to another.
2	Data Link Layer	Responsible for physically passing data from one node to another.
1	Physical Layer	Manages putting data onto the network media and taking the data off.
Packet	A piece of a message transmitted over a packet-switching network. One of the key features of a packet is that it contains the destination address in addition to the data. In Internet Protocol networks, packets are often called <i>datagrams</i> . Also see under packet switching.	
Packet Switching	Refers to protocols in which messages are divided into packets before they are sent. Each packet is then transmitted individually and can even follow different routes to its destination. Once all the packets forming a message arrive at the destination, they are recompiled into the original message.	

PAP	Short for Password Authentication Protocol , the most basic form of authentication, in which a user's name and password are transmitted over a network and compared to a table of name-password pairs. Typically, the passwords stored in the table are encrypted. The Basic Authentication feature built into the Hypertext Transfer Protocol uses PAP. The main weakness of PAP is that both the username and password are transmitted "in the clear"—that is, in an unencrypted form. Contrast with CHAP.
Parallel	Refers to processes that occur simultaneously. Printers and other devices are said to be either <i>parallel</i> or serial. <i>Parallel</i> means the device is capable of receiving more than one bit at a time (that is, it receives several bits <i>in parallel</i>). Most modern printers are parallel.
POP	Short for Post Office Protocol , a protocol used to retrieve e-mail from a mail server. Most e-mail applications (sometimes called an e-mail client) use the POP protocol, although some can use the newer Internet Message Access Protocol (IMAP).
PPP	Short for Point-to-Point Protocol , a method of connecting a computer to the Internet.
PDC	Short for Primary Domain Controller . A PDC maintains the user database for an entire Windows NT domain and its user accounts are said to be domain accounts. When a user tries to log on to an NT domain a request is sent to the PDC which then validates the user.
Proprietary	Privately owned and controlled. In the computer industry, <i>proprietary</i> is the opposite of open. A proprietary design or technique is one that is owned by a company. It also implies that the company has not divulged specifications that would allow other companies to duplicate the product.
Proxy	A server that sits between a client application, such as a Web browser, and a web server. It intercepts all requests to the web server to see if it can fulfill the requests itself. If not, it forwards the request to the web server.
RADIUS	Short for Remote Authentication Dial-In User Service , an authentication and accounting system used by many Internet Service Providers (ISPs). When you dial in to the ISP you must enter your username and password. This information is passed to a RADIUS server, which checks that the information is correct, and then authorizes access to the ISP system.
RAID	Short for Redundant Array of Independent (or Inexpensive) Disks , a category of disk drives that employ two or more drives in combination for fault tolerance and performance. RAID disk drives are used frequently on servers but are not generally necessary for personal computers.
RAS	Short for Remote Access Services , a feature built into Windows NT that enables users to log into an NT-based LAN using a modem, X.25 connection or WAN link.

RFC	Short for Request for Comments , a series of notes about the Internet, started in 1969. An RFC can be submitted by anyone. Eventually, if it gains enough interest, it may evolve into an Internet standard.
	Note: The Internet Architecture Board (http://www.iab.org) is a technical advisory group who manages and publishes RFCs.
Router	A device that connects two LANs. Routers are similar to bridges, but provide additional functionality, such as the ability to filter messages and forward them to different places based on various criteria.
SAM	Short for Security Account Manager . The SAM tracks all information that relates to the account restrictions created on a computer.
Serial	"One by one." Serial data transfer refers to transmitting data one bit at a time. The opposite of serial is parallel, in which several bits are transmitted concurrently.
S-HTTP	Short for Secure HyperText Transfer Protocol . An extension to the HTTP protocol to support sending data securely over the World Wide Web. S-HTTP is designed to send individual messages securely.
SMTP	Short for Simple Mail Transfer Protocol , a protocol for sending e-mail messages between servers. Most e-mail systems that send mail over the Internet use SMTP to send messages from one server to another; the messages can then be retrieved with an e-mail client using either POP or IMAP.
Sniffer	A program and/or device that monitors data traveling over a network. Sniffers can be used both for legitimate network management functions and for stealing information off a network. Unauthorized sniffers can be extremely dangerous to a network's security because they are virtually impossible to detect and can be inserted almost anywhere. This makes them a favorite weapon in the hacker's arsenal.
SNMP	Simple Network Management Protocol was developed so that network administrators could remotely monitor and control network equipment such as bridges, routers and switches.
Spam	Electronic junk mail or junk newsgroup postings. Some people define spam even more generally as any unsolicited e-mail. However, if a long-lost brother finds your e-mail address and sends you a message, this could hardly be called spam, even though it is unsolicited. Real spam is generally e-mail advertising for some product.
SSL	Short for Secure Sockets Layer , a protocol developed by Netscape for transmitting private documents via the Internet. SSL works by using a private key to encrypt data that is transferred over the SSL connection. Web pages that require an SSL connection start with <i>https:</i> instead of <i>http:</i> .

Switches	In networks, a device that filters and forwards packets between LAN segments. Switches operate at the data link layer (Layer 2) of the OSI Reference Model and therefore support any packet protocol. LANs that use switches to join segments are called switched LANs or, in the case of Ethernet networks, switched Ethernet LANs. See also Packet Switching, Packets.
T1	A dedicated phone connection supporting data rates of 1.544 Mbits per second. A T-1 line actually consists of twenty-four individual channels, each of which supports 64 Kbps. Each 64-Kbps channel can be configured to carry voice or data traffic. Most telephone companies allow you to buy just some of these individual channels, known as fractional T-1 access.
TACACS+	Short for <i>Terminal Access Controller Access Control System</i> . The plus indicates that this is a proprietary solution offered by Cisco Systems Inc. It is an authentication, authorization, accounting protocol for systems sold by Cisco.
TCP/IP	Most networks combine IP with a higher-level protocol called <i>Transport Control Protocol</i> (TCP). Different from just IP, TCP/IP establishes a connection between two hosts so that they can send messages back and forth for a period of time.
TFTP	Short for <i>Trivial File Transfer Protocol</i> , a simple form of the File Transfer Protocol (FTP). TFTP uses the User Datagram Protocol (UDP) and provides no security features. It is often used by servers to diskless workstations, X-terminals and routers.
Telnet	A terminal emulation program for TCP/IP networks such as the Internet. The Telnet program runs on your computer and connects your PC to a server on the network. You can then enter commands through the Telnet program and they will be executed as if you were entering them directly on the server console. This enables you to control the server and communicate with other servers on the network. To start a Telnet session, you must log in to a server by entering a valid username and password. Telnet is a common way to remotely control Web servers.
Terminal Server	Terminal servers are stand-alone devices that allow serial devices such as terminals or modems to communicate with the network.
Token Ring	An alternative to Ethernet developed by IBM in the 1970s. In a Token Ring network the computers are schematically arranged in a circle. A <i>token</i> , which is a special bit pattern, travels around the circle. To send a message, a computer catches the token, attaches a message to it, and then lets this combination travel around the network until it reaches the destination computer. The destination computer catches the message-bearing token, retrieves the message and then releases the token back onto the network for reuse.
UDP	Short for <i>User Datagram Protocol</i> , a connectionless protocol that, like TCP, runs on top of IP networks. Unlike TCP/IP, UDP/IP provides very few error recovery services, offering instead a direct way to send and receive datagrams over an IP network. It is used primarily for broadcasting messages over a network.
UNC	Short for <i>Universal Naming Convention</i> or <i>Uniform Naming Convention</i> , a PC

	format for specifying the location of resources on a local area network (LAN). UNC uses the following format: e.g. \\server-name\shared-resource-pathname
UPS	Abbreviation of <i>Uninterruptible Power Supply</i> , a power supply that includes a battery to maintain power in the event of a power outage. Typically, a UPS keeps a computer running for several minutes after a power outage, enabling you to save data that is in random access memory and shut down the computer gracefully.
VPN	Short for <i>Virtual Private Network</i> , a network that is constructed by using public wires to connect nodes. For example, there are a number of systems that enable you to create networks using the Internet as the medium for transporting data. These systems use encryption and other security mechanisms to ensure that only authorized users can access the network and that the data can not be intercepted.
WAN	Short for <i>Wide Area Network</i> . A computer network that spans a relatively large geographical area. Typically, a WAN consists of two or more local area networks.
WINS	Short for <i>Windows Internet Naming Service</i> , a system that determines the Internet Protocol address associated with a particular network computer. This is called <i>name resolution</i> . WINS supports network client and server computers running Windows and can provide name resolution for other computers with special arrangements.
X.25	A popular standard for packet-switching networks. The X.25 standard was developed in 1976.
X.500	An International Organization for Standardization (ISO) and International Telecommunication Union (ITU) standard that defines how global directories should be structured. X.500 directories are hierarchical with different levels for each category of information, such as country, state/province, and city.
xDSL	Refers collectively to all types of <i>Digital Subscriber Lines</i> , the two main categories being ADSL (asynchronous) and SDSL (synchronous). Two other types of xDSL technologies are High-data-rate DSL (HDSL) and Single-line DSL (SDSL). xDSL is similar to Integrated Services Digital Network (ISDN) in as much as both operate over existing copper telephone lines and both require the short runs to a central telephone office (usually less than 20,000 feet). However, xDSL offers much higher speeds.

APPENDIX A

SECURITY CHECKLIST AND SELF-EVALUATION

- Does the organization have a security policy?
- Is the security policy regularly reviewed and updated?
- Are publicly accessible workstations located in areas where they can be easily observed?
- Are vulnerable workstations tied down to prevent removal?
- Are workstations indelibly marked, and is there a record of model and serial numbers?
- Are BIOS passwords secure?
- Are different BIOS passwords used for student and administrative workstations?
- Do workstations have appropriate user profiles?
- Are workstations equipped with antivirus software?
- Is the antivirus software up to date?
- Is there a policy and process for virus-checking diskettes that originate from outside the organization?
- Are student and administrative machines attached to different segments of the network?
- Are passwords and other critical data encrypted before being transmitted on the network?
- Are hubs and switches configured wherever possible to disallow the connection of an unauthorized device?
- Is network equipment located in a secure environment?
- Are servers located in a secure environment?
- Are administrative passwords for servers and network equipment secure, and are they difficult to guess or crack?
- Are administrative passwords changed regularly?
- Are access permissions for shared file systems set correctly?
- Are there appropriate security measures to protect against attacks on server operating systems?
- Are remote access ports protected with passwords or other mechanisms such as dial-back, or caller-ID checking?
- Is activity on remote access ports logged?
- Are the logs of remote access activity regularly reviewed?
- Are administrative systems accessible from public dial-in ports?

- Are any workstations or servers equipped with modems that allow direct dial-in connections?
- Is the connection between the local network and the Internet protected by a firewall?
- Is the firewall configured to provide optimal protection?
- Are firewall logs reviewed regularly?
- Are firewall logs archived?
- Do systems and network administrators keep up to date on current security issues?
- Is a comprehensive security audit carried out regularly?

APPENDIX B

NT WORKSTATION AND SERVER

INTRODUCTION

Many types of information travel over school jurisdiction networks. Some information requires access limitation security to maintain confidentiality or for financial reasons. Other information is intended to be shared freely. The consequences of a security breach may be harm to an individual; financial loss; loss of productivity, information, or files; costs incurred to restore system functionality; simple inconvenience; or no loss at all. Prudent measures do not eliminate these potentials, but they do minimize the possibility of their occurrence.

A well thought out and implemented security policy inflicts minimal or no inconvenience on the system user, while making difficult or impossible unauthorized use, access, or damage. The risks of a security breach should be assessed in the context of each situation and preventative measures taken accordingly. School boards can consider three levels of risk and accompanying measures.

- Level One security is applied to information that must be protected under legislation, or that could result in harm to an individual or financial loss as a result of a security breach. Examples include student information and financial systems.
- Level Two security is applied to information or systems where a breach of security would result in a loss of productivity or require significant effort to restore normal functionality. Information on Level Two systems is not to be considered private by the user. Examples of Level Two security include the jurisdiction's e-mail servers and instructional servers.
- Level Three security is applied to systems where the consequences of a security breach affect only the individual user. An example is peer-to-peer sharing of a teacher's hard drive.

Security involves the protection of hardware, software, and data. Operating systems can not provide physical protection, but operating systems such as Windows NT can provide significant protection for software and data without unduly inconveniencing the user.

Several in-depth references provide information about the security features of the Windows NT (Win NT) architecture. This appendix shows how school jurisdictions might use Win NT. Five scenarios are presented, ranging from a single, stand-alone workstation to a jurisdiction-wide NT domain structure.

OVERVIEW

Windows NT 4.0, introduced in 1996, combines robust multi-threaded architecture with object-by-object access and auditing security, all within an easy-to-use GUI interface. Windows NT is a large suite of "back office" products, including a network operating

system (Windows NT Server), but it also has a "front office" desktop operating system (Windows NT Workstation). In some discussions of Windows NT, only Windows NT Server is being considered, for example, discussions of Win NT versus Novell. In a narrow comparison of Windows NT and other network operating systems, security features are similar. What is significantly different is Windows NT's extension of the majority of the security features of the server to the desktop OS.

In utilizing the security features of NT Workstation, network architects can opt for an end-to-end NT solution or combine the desktop security features of NT Workstation with an alternative network operating system (NOS) such as Novell. The ability of NT Workstation to integrate with an alternate NOS such as Novell allows organizations to leverage their existing investment, continue to take advantage of the features of their existing NOS that may better fit the organization's needs, yet still take advantage of the security of Windows NT Workstation. Whether an organization is using NT or Novell servers, NT security policy templates are available that allow administrators to configure user policies for many mainstream Microsoft applications (such as MS Office or Internet Explorer) and distribute these from the server. Tools are available to create templates for other applications as well.

The advantage of an NT end-to-end solution manifests itself in the seamless integration of the two products (plus the enhanced suite of additional "back office" products) and in the near-identical security configuration procedures between NT Workstation and NT Server. Properly configured, all security policies and features applying to workstation-installed applications and OS can be modified easily (by user or by a group of users) at the server.

As an example, it is possible to use a policy template to defeat access to the configuration dialogs in Internet Explorer. Combined with mandatory user profiles, this feature could be used to force any workstation with Explorer installed to use a school jurisdiction's filtering proxy server when a student logs in. If desired, the same workstation software could allow full configuration dialog access and automatically set profile defaults to bypass the filter when a teacher logs in.

Whether running NT Workstation or NT Server, the casual user can not easily tell the difference between a Windows NT desktop and a Windows 95 desktop. Dig a little deeper though, and dramatic differences become apparent.

From a systems perspective, Windows NT has a much tighter relationship with the hardware it runs on. Unlike Win 95, Windows NT has no ability for software to communicate directly with the hardware. Everything must pass through Windows NT and every action must be sanctioned by the OS. This sanctioning points out one of the major differences between NT and all other common desktop operating systems: integrated security.

Unlike Win 95 or Mac OS, Windows NT requires that users be legitimate (known to the OS). Nobody is allowed to do anything under Windows NT without first identifying themselves through the log-on process. There is no escape key bypass, as there is with Win 95. Of course, to be known to NT, the user must have had an ID created by an administrator. The IDs are created with User Manager and are managed within NT by

the Security Account Manager (SAM), which maintains the database of authorized users.

Log-on security, even for the workstation operating systems, can be established with most NOS platforms to ensure authentication of a user before accessing local or network resources. With Windows NT Workstation, however, full implementation of the security model is much more detailed than the portion provided by the server OS.

NT Workstation (and Server) maintains security information on all named objects and some unnamed objects. To provide this detailed object-by-object security, the volume on the NT Workstation (or Server) must be formatted in NTFS (Windows NT File System). Only then is the Windows NT security system fully available.

The main objective of Windows NT security is to regulate access to all objects including files, folders, printers, and processes that the user can not see. Security information is maintained for each user, group, and object. The features of Windows NT security include User Accounts, User Rights, User Groups, Subjects and Impersonations, Permissions, and in the case of NT Server, Domains. User accounts are SAM database records that contain information about the user, including account name, password, groups, home folder location, profile to be used, and a number of other account privileges. User rights may include setting the system clock or installing device drivers. User Groups allow convenient grouping of user accounts based on similar needs, thus allowing rights, permissions, policies, or subsequent groupings to be accomplished easily. Subjects and Impersonations are used to ensure that a process acting on behalf of a user or other process does not have any different rights or permissions than the user calling the original process has. Permissions (two types) are used to set the access rights on individual objects such as files, folders, printers, or registry keys. Domains are logical groupings of computers (workstations and servers) that share common security and user account information. Domains apply only when NT server is deployed.

Every user account and group has a unique security identifier, commonly known as a SID. This identifier is attached to a security access token at the time of log on and is subsequently provided to every process that operates on behalf of the user. For example, if the user launches Microsoft Word, the user's token is provided to Word. Then, Word can access objects the user is allowed to access but not objects the user is not allowed to access.

The token acts like a key in a lock. Every named object in Windows NT has a security descriptor, part of which is the Access Control List (ACL). The ACL is the lock that the key must fit if the user is to gain access to the object. It also determines the type of access (read-only, change, etc). Configuring the ACL for an object is accomplished with the Permissions dialog as shown in Figure 4. In the Figure, a group of users called MS_Word is being given Read permission to the MS Office directory and all enclosed files and directories.

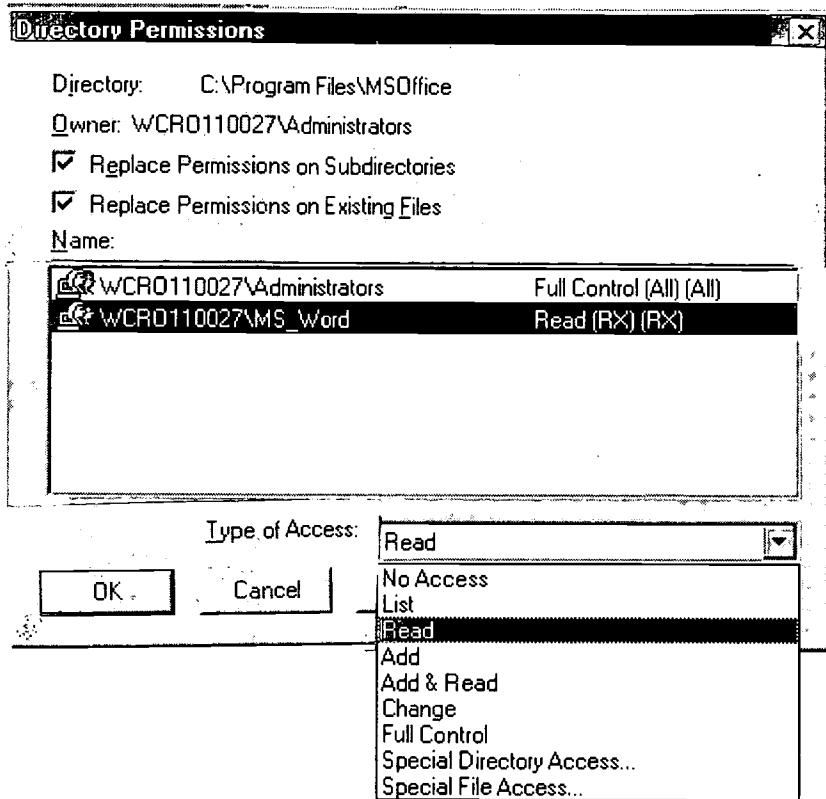


FIGURE 4: THE NTFS PERMISSIONS DIALOG

This object-by-object, detailed level of security known as NTFS permissions is available only if the drive containing the objects is formatted in NTFS format. A second type of permission, share permissions, applies only to shared directories and their contents when accessed over the network.

A key function of any secure OS is the ability to audit. Auditing serves many functions, including logging of successful or unsuccessful access to secured files or directories, troubleshooting network issues, or even determining the source of stress points that are "killing" network performance. NT Workstation and server both have the ability to audit events such as those shown in Figure 5, the Audit Policy Dialog within UserManager.

Once File and Object Access auditing has been enabled, the audit configuration for individual objects must be configured. This configuration is similar to configuring NTFS permissions for the objects, as shown in Figure 6, the Object Auditing Dialog. Using this dialog, administrators choose the access events and the accounts or groups of accounts to audit for a given object. In the example shown, any failure to read an object, or any success or failure to delete an object within the MSOffice directory by any domain user is being audited. These events, as well as all other security audit events, are recorded in the security log.

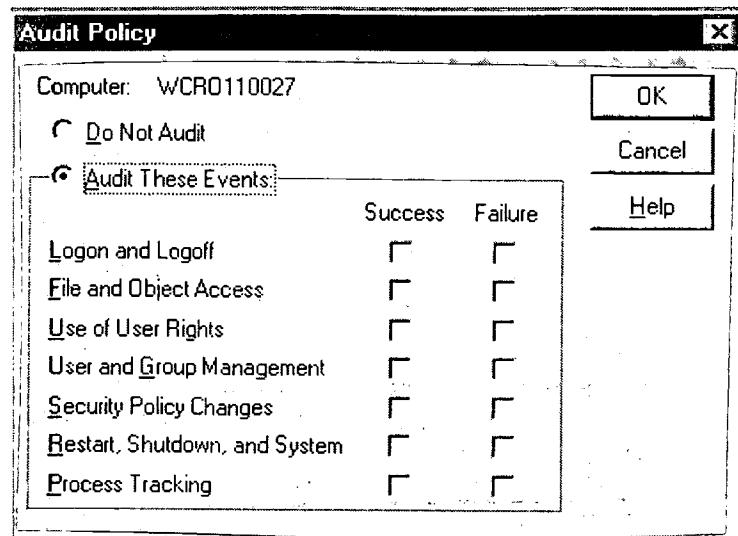


FIGURE 5: THE AUDIT POLICY DIALOG

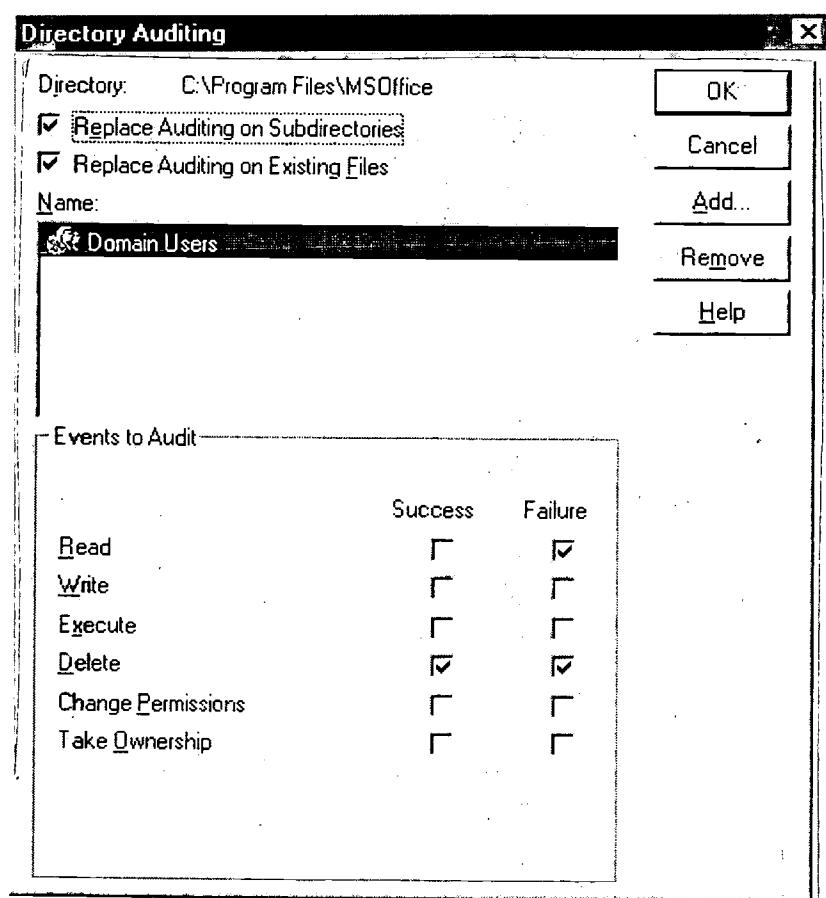


FIGURE 6: THE OBJECT AUDITING DIALOG

STAND-ALONE NT WORKSTATION

While usually thought of as a network-aware OS, in its most basic installation Windows NT could be utilized to secure a stand-alone workstation. Suppose a teacher has a computer in the classroom that has on its hard drive both student assessment files and files for the students to work on.

Remember that nobody gets to use an NT workstation without an ID. One could just leave the administration account with no password, thus allowing anyone to log on and essentially bypass the security of the workstation. But consider the following:

- The system administrator is all-powerful, and it is not that difficult for students to inadvertently or intentionally disrupt the operation of the system by deleting a file, moving a file, or editing the registry.
- Since student assessment files are stored on the computer, the teacher needs to be concerned that a student may tamper with these files.
- A great deal of effort went into developing the lesson plans and template files for student use. Students need access to these files but the teacher does not want them changed, deleted, or otherwise tampered with.
- Some files that board policy requires be kept confidential are also stored on the workstation.
- Students store their work on this computer, and some of them do not always get along well with each other. One student could delete another student's files on purpose, or perhaps unintentionally.
- The teacher across the hall does not have a computer in his/her room and wants to use this one to do some preparation and assessment.
- The school is available for community use and tournaments. Sometimes the room is accessible to visiting students.
- It appears someone used the computer over spring break.

DESIGN CONSIDERATIONS

This seemingly simple and common scenario in non-networked environments actually calls for a fairly detailed configuration. The following configuration could be used to accomplish the requirements detailed above:

- On this and all other workstations, BIOS passwords are enabled to prevent unauthorized changes.
- Boot sequence is set to C:, A:, to impede booting from a DOS floppy disk.
- The Boot.ini timeout for the floppy drive is set to zero. This is a redundant measure similar to the one above.
- If the hard drive is not formatted in NTFS format and the A: drive is enabled, it could be possible to access the files on the hard drive by booting with a standard DOS system floppy disk. Formatting the hard drive as NTFS will prevent access from DOS even if the machine happens to be booted from a standard DOS floppy. (Although shareware programs are available to allow NTFS file read

access when booted from a DOS floppy, this will add an additional hurdle.) NTFS format also will provide the needed object-level NTFS security.

- The Administrator account name is changed to something less easy to guess.
- Keeping visiting students and others out can be accomplished by disabling the guest access account and ensuring the machine is either shut down or always left in a logged out state.
- Users are required to set screen saver passwords.
- Rather than allowing the default user storage location (in C:\WINNT\PROFILES%\USERNAME%) to be used, a documents folder is created for the storage of user files at C:\DATA\USERS where they can be easily located for backup. For most recent applications, the registry can be modified so that the user's folder is the default save location. For others, this must be done within the application. For some lean-featured applications this can not be set at all, so the user must manually navigate to the user's folder.
- A separate folder, C:\DATA\HANDOUTS, is created for teacher-created assignment templates.
- The general guideline for setting NTFS permissions is that nobody has change access to a file or folder unless they need such access. This prevents undesirable configuration changes and provides an extra measure of virus protection.
- Default NTFS security permissions are set using Windows Explorer from the C: drive icon down the tree such that the administrator has full control and users have read-only access. This serves to set the default permissions for all file system objects such that no one can change the system configuration except administrators. From this default setting, only folders and files that the user needs to change permission will be changed. Similarly, to protect the OS from tampering, only portions of the WINNT folder that require change access are modified from the default read-only status.²
- User Manager is used to create unique IDs for the two teachers and for each of the students. A template user ID is created specifying the users' file storage location using the %USERRNAME% variable within the profiles dialog as shown in Figure 7. By copying the template ID when creating the required IDs, home document folders are automatically created with NTFS permissions specifying that only the user has access to the folder.

² The exact file and folder permission details will vary, depending on the applications installed. System administrators may have to experiment with some file, folder, and even registry permissions, especially if the application pre-dates Windows NT 4.0. For example, the permission settings to allow clip art to be used with PowerPoint 95 by all users rather than just by the installing user requires resetting the security on a registry key after installation. Similarly, Corel Office 7 writes two files within the WINNT folder, so change permission is required to the folder the first time WordPerfect is launched, and the two files require change permission for all users thereafter.

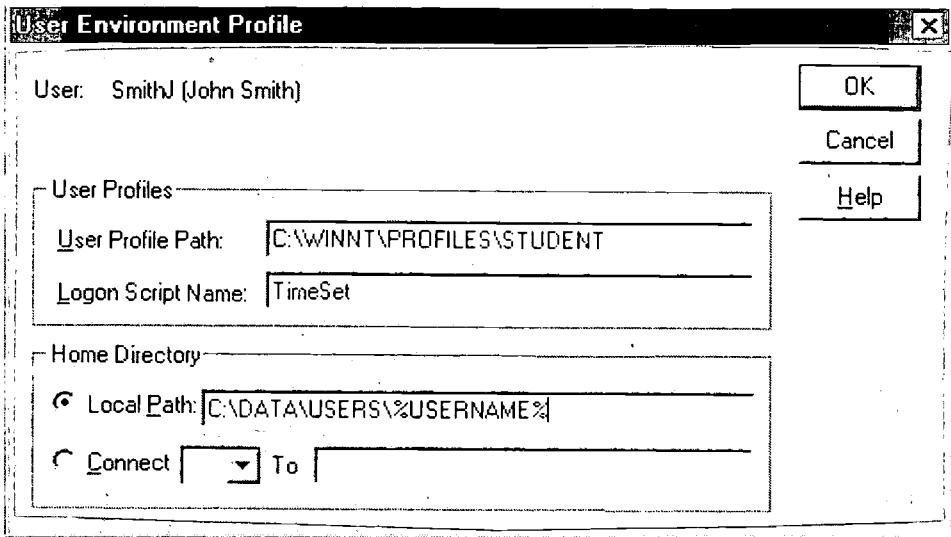


FIGURE 7: THE USER PROFILE DIALOG

- User Manager also is used to create a local group for students as a convenience in managing the student IDs.
- The printer security permissions are set such that students can print but not manage documents.
- NT policy editor is used to edit the registry for each student such that the C: drive is hidden, Windows Explorer is disabled, the Settings and Find items from the Start menu are removed, and a number of other preventative measures (depending on the installed applications) are taken.
- Hiding the C: drive will make it more difficult to gain access to the included games or other applications such as QBasic that are installed on the C: drive. If more assured measures are required, these objects can be specifically set to "no access" for students using Windows NT Explorer. Read access must be maintained to the C:\WINNT folder, however, or the OS will not function for the user.
- Windows NT Explorer is used to set NTFS permissions so that students are only allowed to read, but not change, the templates created for their assignments in the C:\DATA\HANDOUTS folder.
- Windows NT Explorer also is used to set NTFS permissions such that students can not launch the marks application used for assessment. NT's built-in object level security (NTFS Security) has automatically secured access such that users have access only to their own data folders.

The user ID data created when User Manager was used to create user accounts is stored by NT in the SAM database and is utilized by the SAM when a student logs in to create a security access token. The token is used as the key to unlock the objects the student has access to. If a student tries to access another student's files, the token does not fit the lock, and access is denied. If the student tries to access the handouts folder, the read-only lock is opened, and the student can read the files. Even though the C:

drive is hidden in My Computer, it is visible in the Save As... dialog box. If students try to save files in the WINNT folder, they will not be able to since they have read-only access.

The security requirements described above are commonly associated with a network operating system where the user is restricted from accessing certain network resources such as other users' files. With NT workstation, all of the security resources of a network OS and more are essentially built into the workstation OS. In this scenario, NT Workstation performs both client and server functions, operating as if it were a server to itself.

Limited facsimiles of the security provided as an integral component of NT Workstation are available as third party add-ons to Windows 95 or Mac OS but they lack the comprehensiveness of every object having security information and the full integration with the workstation and server OS. Since they are developed by third parties as overlays acting at fairly low levels of the OS, there are often unforeseen and difficult-to-resolve conflicts with the native OS.

NT WORKSTATION IN THE WORKGROUP

Suppose there are now two computers in the room, and a small workgroup network established. The computer security considerations described in the above scenario still apply, but the following also should be considered:

- Students need to use whichever computer is available.
- Student files should be stored on one computer so that the current version of a file is always available.
- The assignment templates also need to be accessed from both computers.
- The printer needs to be shared as well.

DESIGN CONSIDERATIONS

- To share resources, the original computer acts as the workgroup server.
- All of the NTFS permissions continue to apply to the local resources on the workgroup server, regardless of whether the resource is accessed locally or over the network.
- For network access, share-level security now comes into play and complements NTFS security.
- The C:\DATA\USERS folder is shared and the share-level permission is set to change for students and teachers as shown in Figure 8. This will not overwrite the NTFS security on other users' folders, nor the read-only status of the parent USERS folder. When both share permissions and NTFS permissions apply to an object, the most restrictive of the two will apply. In this case, change level share permission ensures that students can change the contents of their own files and folders, but NTFS permissions prevent them from changing the contents of anyone else's.

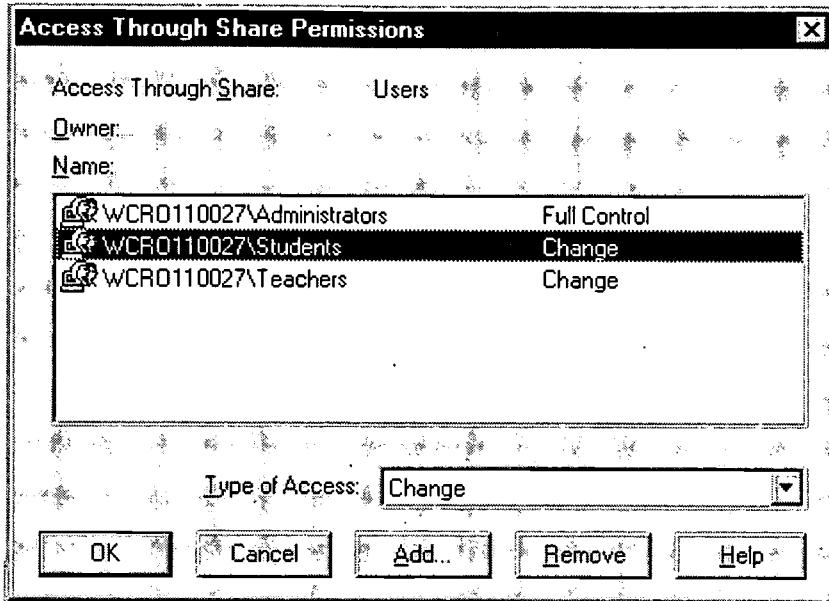


FIGURE 8: THE SHARE PERMISSIONS DIALOG

- The C:\DATA\HANDOUTS folder containing the assignment templates is shared from the original computer, and the share-level permission is set to read-only for students.
- The printer is shared.
- The SAM of the first NT Workstation is unique to it, and it is not shared to the second workstation. Consequently, all of the accounts must be duplicated on the second workstation to allow students to sign on. The template on the original workstation directed NT to the student's home folder on a local drive. On this workstation, the template must direct NT to the shared folder of the original workstation.
- The remaining client workstation configuration is identical to the first one, except that the data folder is not created.

NT has two types of file access security. Share-level security is similar to Win 95 or the sharing security of other operating systems. Share-level security is applied to shared folders, but only when accessed over a network. You can secure network access to a Win NT shared folder in a manner similar to that used to secure network access to a shared folder under Windows 95 or Mac OS. With Win 95 or Mac OS, if you have access to the machine sharing the folder, share security is bypassed. The same is true for Win NT, but remember you must identify yourself to NT before you can access the workstation. In the example above, the files were secured on the Win NT workstation when it was a stand-alone machine. The tool used originally was the other, much more comprehensive kind of file security available under NT called NTFS security. NTFS security is available only if the drive containing the files is formatted using NTFS format, and it is NTFS security that is used to secure access to individual objects.

When files are accessed locally, only NTFS security applies. When files are accessed via a network share, both NTFS and share-level security apply, with the most restrictive one being dominant.

In this scenario, a separate SAM had to be created on each workstation. As the workgroup grows either in number of users or workstations, this becomes extremely awkward and mistake prone. For example, a mismatch in the case of passwords on any one machine in a workgroup will prevent access to the other machines in the workgroup. Similarly, users who change their password on one machine have to change the password on every machine. In fact, in the above scenario, the number of students is likely close to thirty. Creating two IDs each is already burdensome with just two workstations. Besides, NT workstation will allow only ten connected workstations at a time. The solution is to introduce NT server and create an NT Domain.

NT SERVER AND THE SINGLE SERVER DOMAIN

Suppose the above workgroup grew to four computers, or consider a computer lab situation. A cursory look at NT Server will not reveal significant differences between it and NT Workstation, yet the differences are very significant. They include the concept of a domain and the subsequent extension of user, group, and security structures beyond the local computer. Additional network services also are available such as DHCP or Services for Macintosh.

From a security point of view, the big advantage is the ability to centralize the SAM database and apply security measures over the entire domain based on configuration actions at the server. A domain must have a primary domain controller, and this primary domain controller holds the master SAM data base for the domain. It is the one and only place where accounts can be added, changed, or deleted, but it can be accessed from other domain servers (or workstations provided the required administrative tools are installed). Domain controllers also provide authentication services to users when they try to access their accounts.

In this scenario we establish a computer lab facility that provides secure storage of student files, sharing of central resources such as printers, minimal network stress, and still allows the teacher to control which applications are used by which student on a daily basis. Things to consider:

- As the number of workstations or users increases, the need to replicate SAM databases quickly becomes unworkable. In this case there are thirty workstations, and an ID is required for every student in the school.
- Centralizing applications makes it relatively easy to control access, but downloading applications to local workstations can cause a significant amount of network stress, particularly in school situations, where it is common for thirty users to require the same application at exactly the same time.
- All of the same security needs noted above continue to apply, and the additional ability to easily enable and disable applications installed at the workstation is needed.

DESIGN CONSIDERATIONS

All of the above NTFS permissions still apply on the NT server, which will now act as the storage location for the user files. The server is configured in a way similar to the previous workgroup server as far as NTFS permissions and share permissions are concerned. The following enhancements apply to the facility:

- NT Server is configured as a primary domain controller.
- No user accounts are established at the workstation.
- All NT workstation SAMs have only the local administrator account enabled.
- All NT workstations are added to the domain. The log-on dialog will now have a third field added to indicate whether it is the domain or the local workstation that the user is to log on to. Since there is only an account for users in the domain, there is no option for a user to log on without being authenticated by the domain controller.
- User rights at the server are set to ensure that only an administrator can log on at the console.
- All user accounts are created at the domain controller, and consequently apply across the entire domain.
- Home folders are established using UNC paths that direct files to the server share. The share is mapped as a drive on the workstation at log on.
- User profiles for domain users are established as roaming profiles, and student account profiles are directed to use a common mandatory profile. Teachers retain control over their own profile.
- Start menu items such as Run, Settings, Find, and Shut Down are removed from the mandatory profile start menu.
- Default directories for all applications are set to the mapped home folder.³
- Policies are now set globally at the domain controller and Network Neighborhood is removed from the student desktop with the policy editor as shown in Figure 9.

³ Some applications (MS Word 95 for example) use UNC paths rather than mapped drive letters to track the default folder location. In these cases, Network Neighborhood must be retained to ensure the Save As... dialog defaults to the Users share. To accomplish the same objective as hiding Network Neighborhood, the contents of Network Neighborhood can be hidden instead. UNC default directories will then continue to function.

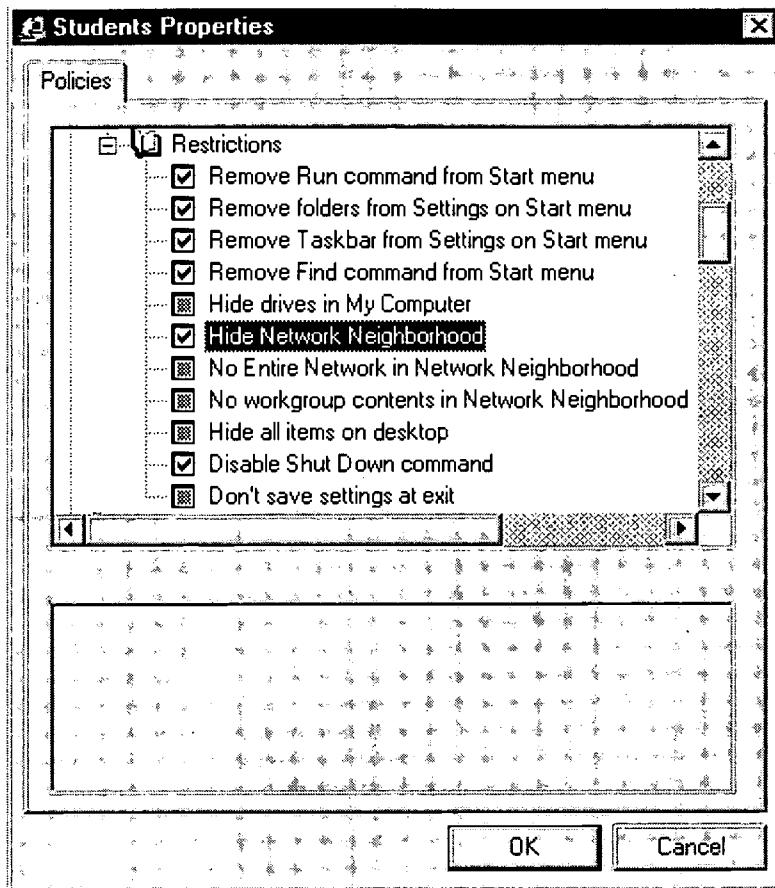


FIGURE 9: EDITING POLICIES FOR THE STUDENT GROUP

- Common core applications such as Microsoft Office are installed on the local workstation hard drive to avoid the network saturation that would occur during simultaneous launch from the server.
- Local groups are established at each workstation for each application.
- Global groups are established at the server for each application.
- NTFS permissions are set at the workstation such that for each application's folder and child objects, the administrator has full control, but the local group for that application has read-only control. All other groups and users are removed from the ACLs.
- The global group is placed in the local group at the workstation.
- Students are moved in and out of the global group to control access to the applications at the workstation.
- A redundant hard drive is included in the configuration, and the two drives are mirrored.
- Depending on the nature of the data and the consequence of loss, a tape backup device also should be considered with data backed up daily and one tape always kept off-site.

This installation allows someone with intermediate administrative privileges to easily enable and disable applications on a per user or per group basis for every workstation by making simple adjustments at the server. Configurations are protected as before, with the additional precaution of eliminating the ability to browse the network. Although this presents little threat in an isolated computer lab network, it provides an additional source of distraction for students.

NT WORKSTATION IN A LAN WITH ADMINISTRATIVE AND INSTRUCTIONAL FUNCTIONS

In this scenario, the school network includes both instructional and school office administration functions. (To avoid confusion with NT administration, we refer to the school administration function as the office function.) Depending on its size, the school may or may not be inclined to dedicate an additional server to the office function. The following should be considered:

- There may be an inclination to assume that NT security features are able to isolate office data from students. This is true, but the quality of a lock is of no consequence if the key is left in it. Suppose someone with office data access privileges neglects to log off at a station accessible to students. The entire security scheme is defeated. Another possibility is that the share or some object NTFS permissions are not correctly set.
- Each function should have its own domain, with no office data stored within the domain where students have accounts.
- NT Server includes all of the functions of NT Workstation. To avoid the cost of a dedicated office server, the server also could function as an NT workstation.
- Since office workstations are generally located in secure areas, storing office data on this computer may be an acceptable compromise. This approach presents far less of a security risk than storing office data on the instructional server.
- It may be desirable to allow office access by office users to instructional resources.

DESIGN CONSIDERATIONS

- If the office workgroup consists of five or more users, consider a dedicated server.
- As in all of the scenarios presented above, change the name of the Administrator account, but use a separate name for each domain administrator account.
- If the cost of a dedicated server is not justified, configure one of the workstations using NT Server. The choice of workstation to act as server should be based on the demands users place on their machines. The least demanding user should become the server.
- Configure the server with the same NTFS permissions as you would a workstation, including local group access control over applications.

- Configure the shares as you would for a dedicated server. Make no changes to the way the person normally using this station accesses his or her home folder. In other words, access to the home folder should be through UNC mapping rather than directly. This will ensure that the users' profiles act the same way no matter which workstation they use.
- User rights will have to be adjusted to allow the user to log on locally at the server.
- To allow office users access to instructional resources, a one-way trust relationship could be established. This would allow office users to have access to instructional resources, but would not allow instructional users to have access to office resources.

NT AND THE ENTERPRISE

In school jurisdictions with large enterprise networks, the scope of security issues is considerably broader. The ability to browse the network now extends beyond the building. But along with increased scope of security issues comes increased fault tolerance capabilities and increased off-site support capabilities. Under NT 4, UNC computer names will not be seen beyond the local network unless a WINS server has been enabled and the browsing computer added to the WINS database. This can be used to a security advantage.

In NT Domains, servers can be configured as either a primary domain controller (PDC), a backup domain controller (BDC), or a member server. Each domain will have only one primary domain controller, although a backup domain controller can be promoted to the primary domain controller. Both PDCs and BDCs handle authentication, but only the primary will allow accounts to be added, changed, or deleted.

A school jurisdiction has a large TCP/IP-based enterprise network with all instructional and administrative workstations connected. NT Workstation and Server are the only operating systems used on office workstation computers. School offices vary in size from two workstations to nine. Any office with more than four workstations also has a dedicated NT server. Those with four or less use NT server on one of the workstations, and it acts as a file and print server to the other workstations. Instructional servers are all dedicated NT servers, but workstations use one of NT Workstation, Windows 95, or Mac OS. Consider the following:

- NT servers providing Macintosh services utilize AppleTalk and will therefore be visible in the chooser of Mac OS workstations.
- WINS services are required before a computer outside the local area network becomes visible in the network browser (or Find Computer command). Once WINS is enabled, NT Servers, Workstations, or Win 95 stations must still be configured before they can see beyond their LAN.
- A secondary WINS server can be configured for each computer so that if the primary does not resolve the address, the secondary will be consulted.
- If the IP address of a workstation is known, it can be accessed via the run command regardless of WINS configuration.

- Domains do not normally share their resources with users from other domains.
- If an ID and password are identical in two domains, the permissions and rights of each domain will be available to the user. This is similar to the workgroup in the NT Workstation in the Workgroup scenario.
- One-way trust relationships can be established such that users in one domain can have access to the resources of another, but not the other way around.
- Administration and support is enhanced if support personnel can access workstations and servers remotely.

DESIGN CONSIDERATIONS

- Global groups are established for all students in each school, and as many additional groups as are needed to conveniently manage student accounts. Examples might include a group for all students in each class or grade.
- NTFS and share permissions are configured as above, except that printing is set so that only students in a given school can use a printer in that school.
- If desired, user rights to workstations may be set so that only users in the local school can use a workstation.
- A separate school office domain is established. A single school administration domain will provide a BDC in every school for local user authentication, but since these servers are all part of a single domain, each provides fault tolerance of the SAM database for each other.
- Similarly for the instructional domain, enterprise-wide domains are established to provide fault tolerance for the SAM. The number of instructional domains would ideally be one, but the size of the SAM (number of users) and the availability of WAN bandwidth may dictate that two or more domains be established to ensure reasonable performance when account changes are made. Remember that authentication can happen at any domain controller, but changes to the SAM require access to the PDC.
- The central office should have its own domain, and its own administrator account name.
- The technology support group should have its own domain and administrator account name.
- WINS configuration on student workstations should be left blank so that these stations can not see beyond their local network. This provides a bit of redundancy as the Find command and Network Neighborhood have already been removed with policies.
- If the central office is completely on one LAN (in one building), then WINS is not used on the servers to ensure that these servers are not seen outside the building. If visibility is required because the central office functions out of two or more buildings, then WINS configuration on central office servers is set so that they only see each other. This is accomplished by enabling WINS on one central office server and having the other(s) configured to use it.

- School-based servers and school office workstations all use a separate WINS server somewhere in the school jurisdiction.
- Technology support workstations are configured to use both WINS servers.
- Trust relationships are set using one way trusts only. The instructional domain trusts the school office domain and the technology support domain. The school office domain trusts the technology support domain. The central office domain trusts the technology support domain. This arrangement allows the technology support domain to gain access to all domains for support services and act as a Master Domain. (User files are still protected through NTFS permissions.) School office users can access the resources of the instructional domain. All other users are confined to the resources of their own domain. Since student stations do not have WINS configured, they can not access resources such as printers outside their own network. The example in Figure 10 shows that the WC_INST domain trusts the members of the WC_TECH domain, but the members of the WC_INST domain are trusted by no other domains.

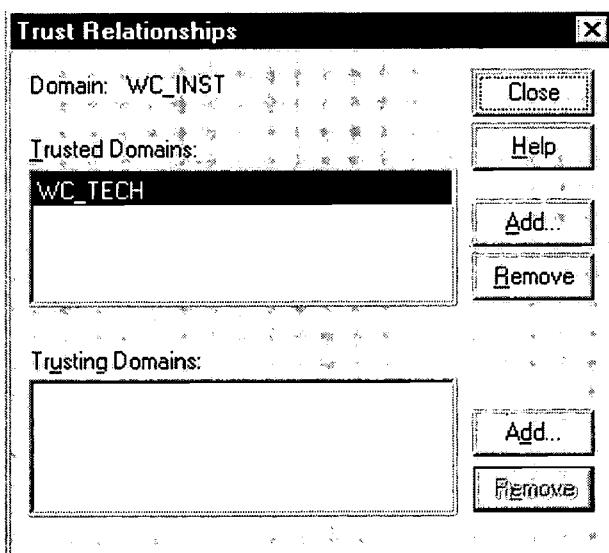


FIGURE 10: THE TRUST RELATIONSHIP DIALOG

- In situations where school networks are bridged rather than routed, access to resources will have to be limited with share and NTFS permissions.
- Services for Macintosh are configured only on instructional servers as there are no Macintosh computers in offices.
- AppleTalk is disabled over the wide area to avoid accessing AppleTalk services such as network printers in other schools. This also is a good idea to avoid having the frequent broadcasts of AppleTalk consume limited WAN bandwidth.
- Each domain has at least a PDC and a BDC to provide SAM redundancy. In addition, each server is equipped with a mirrored hard drive.
- At a minimum, school office, central office, and technology support servers have tape backup facilities that are used daily including keeping an off-site backup.

- Messenger services are configured such that technology services is automatically informed of critical events.

CONCLUSIONS

Information systems security is often focused on controlling access to information and resources. Only those who are entitled to access are able to obtain it. A security issue that is sometimes overlooked is the financial or functional loss when effective control is not applied to total cost of ownership. The application of standards is an effective tool to control cost of ownership.

"Technology plans are doubly cursed when staff needs are underestimated and a standards policy is not articulated. No step can do more to control the proliferation of new staffers than an effective set of standards. Jurisdictions that choose to adopt a freewheeling systems policy will see their staffing requirements grow uncontrollably as their projects increase in size. Such jurisdictions will then have to choose between two equally unattractive options: either to understaff, or to pay the price, no matter how high. Most jurisdictions will choose to understaff, thereby hamstringing their system far into the future" (Weiss, 1996, p. 409).

In addition to security of information, NT provides security of cost of ownership, either through NT Workstation and other network operating systems, or through an end-to-end NT solution. Effectively implemented, Windows NT can increase system reliability and affordability by protecting standard configurations while still providing teachers with the tools to control student access to resources. Increasing reliability considerably reduces the demand for support resources and therefore the total cost of ownership.

Technology infrastructures should be characterized by adjectives such as reliable, affordable, supportable, and appropriate. A good security implementation will protect access to information and resources but not impede legitimate users from easily accomplishing their tasks.

REFERENCES

Brown, Todd C. and Norton, Peter. (1997). *Peter Norton's Maximizing Windows NT Server 4*. Indianapolis, Indiana: SAMS Publishing.

Garms, Jason. (1996). *Windows NT 4 Server Unleashed*. Indianapolis, Indiana: SAMS Publishing.

Minasi, Mark. (1998). *Mastering Windows NT Server 4*. Fifth Edition, San Francisco, California: Sybex Incorporated.

Minasi, Mark and Campbell, Patrick T. (1996). *Mastering Windows NT Workstation 4*. San Francisco, California: Sybex Incorporated.

Moore, Sonia Marie, et al. (1996). *Microsoft Windows NT Server Resource Kit*. Redmond, Washington: Microsoft Press.

Moore, Sonia Marie, et al. (1996). *Microsoft Windows NT Workstation Resource Kit*. Redmond, Washington: Microsoft Press.

Sanna, Paul. (1996). *Special Edition Using Windows NT Workstation 4.0*. Indianapolis, Indiana: Que Corporation.

Weiss, Andrew M. (1996). "System 2000—If You Build It, Can You Manage It?" *Phi Delta Kappan* 77:6, pp. 408–415.

APPENDIX C

NOVELL NETWARE

INTRODUCTION

Novell NetWare was an early innovator in developing a network operating system for Intel-based systems. The first generations of NetWare were based on a simple design of a single file server with multiple workstations. The NetWare focus was to provide file and print services to multiple machines by allowing network resources to appear as extensions of the local machine.

The original NetWare design was to have all network information stored in a server file called *the bindery*. In this original design model, users authenticated to a server to use its resources. The security within this model allowed the administrator to restrict a user's access to any network resource and delegate authority to manage network objects. While the information in the bindery was encrypted and protected from even the administrator's view, IDs and passwords were passed as clear text across the network wire.

In large networks with multiple servers, users were required to authenticate to each server in order to use the network resources. This created an additional administrative burden as users had to be defined on every server. This design model did not work well in large multiple server networks because of this increased administrative load and the difficulties in presenting a seamless integrated structure to the end user.

In the early 1990s, NetWare recognized the problems with its single server model. Recognizing the advantages of utilizing a centralized database of all network users and objects, Novell developed its own centralized database called Novell Directory Services or NDS.

Novell Directory Services (NDS) is a centralized distributed database of all the objects (users, servers, printers, etc.) that comprise the network. In essence, NDS is a combination of the features found in OSI X.500, Banyan Street Talk, and Novell's ideals of what a large WAN directory should be.⁴ NDS differs from NetWare's older bindery technology in that users no longer authenticate to a server, but to the NDS database instead. NDS contains the database or rules that control interaction between the objects that comprise the network. Because users are no longer tied to a specific server, NDS provides a scaleable and workable solution for multiple server networks.

With the advent of new technologies and an increased market desire for complete network solutions, Novell NetWare is under increasing market pressure to evolve from a Network Operating System (NOS) to a complete network environment. Despite this evolution, Novell continues to build on its basic security structure, Novell Directory Services (NDS). Whether NetWare is providing file and print services, providing Internet web publishing, or acting as a firewall, NDS is at work in the background to provide the backbone of the Novell Security Model.

⁴ David James Clark, *CNE Study Guide for NetWare 4.1*, Novell Press, 1995.

As in any network, security starts at a physical, not a logical, level. Physical security involves the users' hands-on access to workstations and the administrators' access to critical components such as file servers and the network backbone such as hubs, switches, and routers.

Physical security is extremely important when it comes to locking down the NetWare server. The NetWare server console (console) is by its nature a security threat. The console is **not** automatically protected: the administrator must lock the console using the MONITOR utility. The console can then be unlocked using either the Admin (top level) password or the password that was supplied when the console was locked.

The console is one of the weakest points in the NetWare security model. A user or administrator who has access to the console has the potential to do almost anything. From the console, a user can access user data without restriction, change network configurations, destroy complete volumes of information, and more. The user is limited only by imagination in ways to utilize the abundant tools available on the server console. The console is usually protected by a "shared" password, and all persons who must have access to the console have the same level of access regardless of the activity they need to perform. Further, activity at the server console is not audited in any useful manner.

By comparison with the server console, once a user has an interaction with a workstation, there is now a logical interaction with the NOS as well as a physical interaction. This interaction with the NOS brings an additional layer of security that is not present at the server console. Because the workstation is accessed by ordinary users, Logical Security is used to protect the network. Logical Security covers such things as a user's ID and password, log-in restrictions, access to the file system, and a user's access to other network objects such as printers, servers, or even other users.

Logical Security is the primary way the Novell security model deals with potential intruders. These potential intruders can be grouped in two classes: internal intruders (or threats) and external intruders (or threats). Without question, internal threats are the greatest risk to a network because they have physical access to the network on a daily basis. External intruders become a threat only if the network is accessible via the Internet, or via a dial-in service. A network isolated from the outside world can not be penetrated unless the intruder has some means of gaining access to either a workstation, a server console, or a segment of the network backbone.

Internal threats can range from a disgruntled employee actively seeking to gain confidential data or attempting to disrupt network activity, to the curious power user testing the limits of the system, or the inexperienced user poking around to find his or her files. While organizations will always have the willfully destructive, network browser utilities such as Network Neighborhood have made it extremely easy for users to get into file systems that were previously accessible only to users familiar with network architecture and the command line utilities (such as MAP) needed to access those file systems. Without question, these graphically based utilities require the network administrator to be particularly vigilant when designing large networks with multiple file systems. Without due care and attention, user curiosity can compromise confidentiality.

External intruders are generally malicious. When a network is connected to the Internet or some other public external network, some sort of firewall must be used. The firewall may be another Novell server running Border Services or another non-NetWare firewall. Without a firewall, the local network becomes a physical extension of the public network.

Another often overlooked threat is dial-in access to the local network. As with the Internet, once a remote access server is in place and connected to the telephone line, the local network becomes a physical extension of a public network. In this case, it is an extension of the Public Switched Telephone Network (PSTN).

Whether the threat is internal or external, Novell has the means and ways of providing security services (logical security) to protect and secure the network.

NOVELL SECURITY MODEL

Before the security that NDS provides can be implemented effectively, some thought must go into the design of a network's security policy. Novell uses a model that looks at Environment, Risk, and Assurance (the ERA Model) to help design and implement a workable security policy.⁵

The last four words of the preceding paragraph should be read again with some thought: "a workable security policy." A policy that is only on paper is worthless. A security policy must be practical, enforceable, and to a large part invisible to the end users. When a security policy restricts users' functional access to the network or becomes annoyingly intrusive, users will find ways to work around it. Users must be aware of the purpose of the security policy, and must agree to support it.

The first step in the ERA Model is analysis of the physical environment. This involves all levels of physical security. As noted previously, access to Novell File Servers must be restricted. Other environmental factors include the network backbone itself. The administrator must look closely at the hubs, routers, bridges, and switches that connect the network. If a potential intruder has access to the physical backbone of the network, physical security has been compromised.

The target audience (those that the network serves) can and should influence the network design and topology, which in turn has an impact on the physical security of the network. Organizations that serve large transient populations (for example, universities and schools) face different security concerns than small businesses. Organizations that have mobile users who dial into the network or networks connected to an external, untrusted network such as the Internet or PSTN have other concerns as well.

The informational assets being protected also need to be considered. The administrator needs to consider how this information will be secured, and how access to this data will be restricted. It is not unheard of to have confidential data backed up to tape, and then have that tape stored in a public facility. Once data is archived, it is no longer protected by the NOS. Data on tape must be protected by physical security or some sort of encryption within the backup process.

⁵ *Novell Research App Notes*, November/December 1997.

The second step in the ERA model is risk assessment. These are the questions to be answered at this stage: "How much can we afford to lose?" "What are the implications of information being made public?" "What is acceptable downtime?" As children we learn, "What goes up eventually comes down." The same is true of network components: all hardware is subject to failure. Ideally, all network components are mission critical, but unfortunately this is not always practical.

The third step in the ERA Model is assurance. "Where do you feel comfortable?" "Once the environment and the risks have been assessed, how far do you go?" There is no set answer. Budgets often get in the way of an ideal security policy. While logical security is fairly cost effective to implement, physical security brings additional costs, especially if it requires redundant hardware or the building of secure backbones (switches as opposed to hubs). When determining whether these additional costs are justified, school jurisdictions must consider that there also is a cost to not having the environment that meets their needs.

An analysis of the ERA Model reveals that threats to network security fall into three basic categories:

1. Hardware failures and other natural disasters
2. Human error or unintentional damage by users
3. Deliberate and willful destruction or sabotage, either internally or externally (See Figure 11.)

The security policy must cover physical security implemented within the network environment and logical security implemented through NDS.

Human Errors	55%
Physical Security Problems	20%
Dishonest Employees	10%
Disgruntled Employees	9%
Viruses	4%
Outside Attacks	1-3%

Source: **NetWare Connection Magazine**, January 1996

FIGURE 11: NETWORK SECURITY THREATS

How logical security is implemented in NDS is a vast subject well beyond the scope of this paper. This paper focuses on the general concepts of how Novell NetWare uses NDS to secure the network. For a detailed study of implementing NDS, see the list of articles provided in the "References" section of this appendix.

Log-in security is the first line of defence in guarding the network. Users must provide a log-in ID and a password to gain access to the network. Log-in restrictions can limit

users' access in a variety of ways; for example, when they log in (time restrictions) or where they log in (address restrictions).

Once users have been authenticated to NDS (logged in to the network), NDS determines which file systems they can or can not access (file system security) and which network objects such as print servers and printers (or another object) they are allowed to interact with. In the world of NDS, printers, print servers, file servers, file systems, users, groups, and all other resources are considered NDS objects. Because of this, NDS provides not only file system security but object security as well. In the same way that file system security determines what files a user may access, create, modify, or delete, object security determines what objects a user may access, create, modify, or delete.

While NDS may appear complex at first, it is in reality a powerful and easy-to-use way of controlling user access. NDS provides a logical and distributed means of linking network users with the network resources they need, and a way of restricting users from the network resources they do not need.

NDS MANAGEMENT MODEL

By its own design, NDS provides a logical way to distribute the administrative tasks within the network. Object rights can be used to give a user the power and rights to administer the portion of the network s/he is responsible for, while restricting access to portions that others are responsible for.

The underlying principle behind NDS is Inherited Rights: rights given at a higher level of the NDS tree flow down the branches to objects. This concept of inherited rights is a valuable tool for network administrators, but in large networks it can be a potential security risk. To block the natural flow of rights, NDS uses an Inherited Rights Filter (IRF) to secure lower portions of the tree.

Jeffrey F. Hughes and Blair W. Thomas provide a complete discussion of designing an NDS tree. (See "References" section of this appendix or

[<http://www.novell.com/nwc/jun.96/ndstre66/>.](http://www.novell.com/nwc/jun.96/ndstre66/)

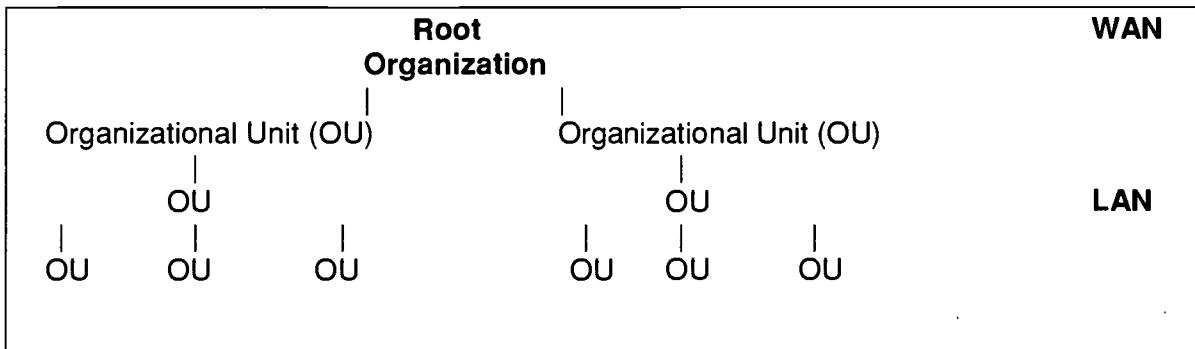


FIGURE 12: A GRAPHICAL REPRESENTATION OF THE NDS TREE

The basic design behind NDS is to group users and the resources they most often use into an organizational unit (OU). These OUs can then be organized to represent an organization's physical layout or departmental structure.

A well-designed NDS tree is essential to maintaining and designing an effective security policy with a Novell Network. The tree provides the means to link users to resources they need and to secure resources from unauthorized access or use.

CLIENT MANAGEMENT AND SECURITY

While the logical security of the NOS is normally the first thought that comes to mind when considering security, that fact that users will be interacting with the network from a workstation must not be forgotten. A useable NOS must integrate with the operating system of the workstation in a way that is a logical and useful extension of the local machine. Users who can find the tools and data they need to do their job are less likely to go exploring.

With DOS-based clients, Novell presented the network as additional drive letters and additional printer ports. It was very easy to shield the user from the network, and present network resources as mere extensions of the local machine.

With the advent of Windows 95 and Windows NT 4.0 and the EXPLORER interface, users become very aware of what are network resources and what are local devices. This change has made the network a more user-friendly environment, but it also has made it easier for users to find data and resources they should not be accessing.

Fortunately, the network administrator has many tools available to help distribute network resources and secure the local workstation. In Windows 95 and Windows NT, the NetWare client works in conjunction with *policies* and *profiles* to secure the workstation. For distributing network applications, Novell utilizes a tool called the NetWare Application Launcher.

It may be desirable to secure the local workstation so that end users can not make changes and restrict how users interact with the network. Consequently, Microsoft has introduced a concept called *profiles* and *policies*. *Profiles* define how users see their local machines and can restrict such things as screen savers, colour settings, and other local environmental settings. *Policies* restrict how users view their network environment, and which aspects of their *profiles* they are allowed to change. To make full use of profiles and policies, the network administrator must be familiar with Microsoft Networking and security within a Microsoft network.

Since the Novell client is now closely linked to the Microsoft networking model, an effective network administrator must now have not only a complete understanding of Novell security but of Windows security as well. While this adds a level of complexity, it also adds flexibility. Having a Novell network is not an either/or approach. In fact, a Novell network can fully integrate and co-operate with a Windows NT network.

A workstation can be a member of a Novell network and a Microsoft network without conflict. This represents a new dimension in network connectivity. Users can now access resources on either Novell or Microsoft networks. The network can be designed so that, while users may be aware that they are accessing a network resource, the resource appears the same whether it is provided by Novell NetWare or Microsoft networking.

NetWare has one tool that simplifies user access to network applications. The NetWare Application Launcher (NAL) used in conjunction with the power of NDS can make applications become network objects in the same way as files, printers, and other traditional network objects. Because these applications are now objects, NDS can restrict who has access to the applications and present the applications the user is allowed to see on the user's workstation desktop. The power of this tool is significant. Instead of installing an application on thirty, sixty, or more workstations, the application is installed only once and the power of NDS is used to distribute that application to all users.

The power of the NetWare Application Launcher is enhanced by another utility called SnapShot. When an application is installed, it often copies files to the local workstation, makes changes to INI files, or makes changes to the Window's registry. SnapShot allows the network administrator to record all of the changes that take place on the local workstation and then record the changes in a template file. This template file can then be used by NAL to make the necessary changes to the local workstation when running an application.

Since NAL, a free utility, is only in its second release, it will be revised and improved. SnapShot is also a free first-generation module. These two utilities are some of NetWare's most significant enhancement releases to date.

NOVELL REMOTE MANAGEMENT UTILITIES

Novell provides simple but effective utilities for remote management of servers and network objects. Remote management consists primarily of two utilities: the NetWare Administrator for managing NDS (objects and file systems) and the RCONSOLE utility for managing file servers.

The NetWare Administrator is a graphical-based utility that allows the network administrator the ability to manage NDS objects. Since everything from users to printers to file systems and applications are NDS objects, only one tool is needed to manage the network. This easy-to-use tool can be used to manipulate any object anywhere in the NDS tree, or with the newest version any object on any NDS tree on the physical network.

The second tool for remote management, the RCONSOLE tool, is a utility that recreates the server console on any client workstation on the network. With this tool, the network administrator can control any server regardless of its physical location. On large geographically diverse wide area networks, RCONSOLE allows an administrator to work directly and easily on a server even if it is hundreds or thousands of kilometres away.

While RCONSOLE and the NWADMIN tools are the mainstay of the remote management tools, there is one extra tool available to administrators running TCP/IP-based networks. The XCONSOLE utility can be used to provide Telnet access to any Novell server that is using TCP/IP. With this tool, an administrator can manage a Novell server from any workstation on the network that can act as an Internet host, regardless of its operating system.

NOVELL AND INTERNAL SECURITY RISKS

Novell has always dealt effectively with internal security risks. With DOS-based clients, it was easy for the network administrator to have the network appear as little more than an extension of the local machine. While the average "power user" had no problem distinguishing between network and local resources, the average user was unable to access resources other than those explicitly assigned. This, however, has changed dramatically with the advent of graphical-based user interfaces and the use of the networking model employed by the newer Microsoft operating systems such as Windows 95/98 and NT.

The user running Windows 95/98 or NT can use Network Neighborhood to which resources are local and which are provided by the network. The user also can actively browse the network to easily find all network resources, including those that the administrator may not want the user to access. This is generally a minor problem in a small network environment, but it can become a major problem in a large environment.

In a large, fast-growing environment, network management often is handled by multiple administrators of varying skill levels. This is where the Network Neighborhood problem is at its worst. The inexperienced administrator tends to assign the users more rights than they actually need. Administrators often solve problems arising from NetWare rights restrictions by removing all restrictions or making the user with the problems *supervisor equivalent*. These short-term solutions often are left in place for long periods of time.

When more than one person administers the same network, it becomes difficult to keep track of who has rights to what and who has administrator privilege. Accurate documentation must be created at the time the network is installed, and those documents must be kept current and up to date. If user rights are not tracked and documented from day one, this problem will eventually become critical.

In the days of NetWare 3.1x, a small utility called SECURITY listed users that were supervisor equivalent and pointed out possible security breaches within the system. This tool does not exist in NetWare 4.1x, so to audit user rights the administrator must look at every NDS object that comprises the network. In a large network with thousands of objects, this becomes an unenviable task. At present, there are two third-party utilities that can scan a Novell network and generate an audit of user rights: the Kane Security Analyst and AuditWare from Cheyenne Software.

NOVELL AND EXTERNAL RISKS

Novell is just now venturing into the world of securing the network from external security risks. By linking to its powerful NDS database, Novell Border Services can provide firewall protection to the network, and provide rules based on individual users as opposed to the network as a whole. Traditional firewalls are not aware of the network objects or users behind them. By linking to the NDS database, Border Services can apply rules to any network object or user.

In the past, rules of a firewall pertained to how the Internet and the internal network were allowed to interact. Therefore, it was very difficult for a traditional firewall to restrict individual users. The firewall saw the Internet, the internal network, and the individual host machines that comprised the internal network. Users were simply not a part of the equation. Border Services has changed all that. By harnessing the power of the NDS database, it can create firewall rules that apply to individual users as well as host machines.

The power of this new technology is staggering for organizations that need to restrict or limit users' access to the Internet. This technology allows an organization to restrict access to Internet services (web browsing, Internet relay chat, etc.) for some users while allowing unrestricted access for others.

While Border Services is a relatively immature product, it is certainly one of Novell's most exciting new technologies.

ADDRESSING HARDWARE RELIABILITY

Hardware failure is the biggest risk that any network faces, and it is an indisputable fact that all hardware eventually fails. Even though hardware has become more reliable in past years, it is a wise precaution to utilize the features built in to NetWare to minimize the risk of hardware failure.

For mission-critical services, Novell offers its SFT-III (System Fault Tolerant) version of NetWare. This technology allows for complete redundant servers. If one system fails in any way, the redundant system provides seamless and transparent protection.

While complete redundancy is ideal, Novell offers drive mirroring and duplexing to deal with that other more common and annoying problem—hard drive failure. Mirroring hard drives and duplexing (mirroring not only the drive, but also the controller) is an extremely cost-effective and easy-to-implement way of securing data and enhancing server reliability.

While data security is important, in a large network the user database (NDS in this case) also is very important because of the time required to recreate it in the event of a disaster. In a multiple server network, NDS takes care of this problem by its very design. No single server is the repository of all network information. Since NDS is distributed across all servers, the database is never at risk from a server failure. If NDS

data is properly replicated, half of the servers could fail simultaneously and the NDS database would still be recoverable.

SUMMARY

Ultimately, network security is not a function of the Network Operating System (NOS) but of the administrator who designs and implements the security. Most NOS's can be secured to address most security concerns. NetWare is no exception.

By design, NetWare provides all of the tools necessary to create a secure networking environment. It is up to the administrator to use those tools to meet the needs of the organization and provide the appropriate level of security.

UTILIZING NOVELL IN PUBLIC EDUCATION

Networks in school environments often pose unique challenges for the network administrator. These challenges must be recognized and appropriate strategies developed to ensure that network security is maintained while allowing a reasonable level of access.

Factors to consider:

- Roaming users: one person using multiple machines
- Student access versus teacher access
- Support staff access versus teacher access
- Administration team access (principal, etc.) versus teacher access
- Student record system
- Curious users/hackers
- Multiple administrators (large networks)
- "After-hours" access (i.e., libraries, night classes, etc.)

In a school environment, careful planning is the key to designing a secure network. By utilizing file permissions on the file server, Novell provides a means to provide or restrict access to specific files or file locations. Because of the design of NetWare file rights (permission flow down), it often makes sense to separate student-accessed files and teacher-accessed files into unique and separate directory structures. By its nature, NetWare is a restrictive environment, so if the administrator maintains a policy of assigning the least rights needed, the network file system will be relatively secure.

The greater challenge is securing the local workstation. In a DOS or Windows 95 environment, the user has complete control and access to the local workstation. Because the user can make changes or modifications to the local operating system, reliability is often compromised. Computers do not work with improper configurations. More frightening is the possibility of unauthorized programs being installed. While an unauthorized program may be as benign as a simple game, it could be as destructive as

a virus, or as security-threatening as a program that captures log-in IDs and passwords. In a volatile environment such as public education, a more secure local operating system is needed.

The solution to the problem of securing the local workstation is utilizing Windows NT as opposed to Windows 95 or DOS. Windows NT is a secure operating system. Users can not access the local PC until they have been authenticated to the local security database. Users can not bypass the log-in screen by pressing escape. If you do not authenticate, you do not access. To work in conjunction with the local security database, the Novell client either dynamically creates the users in the local database or uses a synchronization utility to ensure that the user and password exist in both the Novell NDS and the Windows NT security database.

Windows NT also provides for placing network-type security permission on the local workstation's hard drive. This allows the administrator to restrict access to prevent unauthorized file deletions or additions. This additional level of security should be explored and is certainly a "must" for student-accessed workstations.

Windows NT also allows for the use of *profiles* and *policies*. These tools are provided in the Microsoft Networking Model as an additional means of securing the local workstation. Policies control what settings a user is allowed to change. Profiles control the user's desktop and preferences. Using these tools, the administrator can allow one user to customize his/her environment, while restricting the environment of another. The network administrator would probably want to utilize these tools to secure the workstation so that students are fairly restricted, while allowing some flexibility for teachers and administrators.

While there is no one "recipe" for securing the local workstation, these recommendations should be considered:

- Windows NT environments are easier to secure than Windows 95 and MS DOS.
- Utilize NTFS (file system) restrictions to secure the local hard drive.
- Utilize policies to secure the local workstation's network and system configuration.
- Utilize policies to restrict student access while allowing more control for other users.
- Utilize policies to allow "power users" to control their environment (if appropriate).
- Utilize shared roaming profiles (read-only access) to provide a consistent environment for all students.
- Utilize individual roaming profiles to provide a personalized environment for teachers, administrators, and support staff.

REFERENCES

"Devising an Information Security Policy: Environment Risk, and Assurance." *Novell Research App Notes*. November/December 1997.

Hamilton, Carey and Broyer, Linda. "Netware Security: Making Your Network a Fortress." *Netware Connection Magazine*. January/February 1996.
<http://www.novell.com/nwc/jan-feb.96/security/security.html>

Hughes, Jeffrey, T. and Thomas, Blair W. "Designing an NDS Tree: Three Objectives To Help You Create an Effective Tree." *NetWare Connection Magazine*. June 1996.
<http://www.novell.com/nwc/jun.96/ndstre66/>

"Learning and Applying the Rules of NDS Security." *Novell Research App Notes*. August 1997.

"Network Security for the 21st Century: Concepts and Issues." *Novell Research App Notes*. November/December 1997.

"Overview of the NetWare Enhanced Security Architecture and Configuration." *Novell Research App Notes*. November/December 1997.

APPENDIX D

APPLE MACINTOSH

INTRODUCTION

Macintosh hardware and software have evolved into a suitable computer and operating system for use in schools. The interface has not had major visible changes over fourteen years. While some suggest that this is a disadvantage, it has advantages in time savings over learning a new interface, and working with it. Techniques learned with older Macintosh operating systems can be used with the newer versions.

The original Macintosh was shipped with built-in LocalTalk networking, but now they are equipped with built-in Ethernet networking. If necessary, the Macintosh can still broadcast LocalTalk for communication with older equipment. Thus, in a school with a mixture of older and newer technology, all computers can still be utilized. Macintosh Plus, SE, and Classic computers are still in regular use.

A unique feature of the Macintosh platform is the basic design of the operating system. Since there is no underlying text-based interface, the Macintosh is 100 per cent graphical. There is no conversion done behind the scenes to run older software and there is no overhead created.

A significant feature of Macintosh is security when properly defined. Student access can be limited on the local hard drive (read/write/delete) and on the network.

The operating system allows the Macintosh to access any network resources, and in fact act as a server itself should the need arise. Apple's Network Administrator Toolkit is an "add-on" package that makes the Macintosh a fully networked system from the workstation to the server.

On some platforms, server software differs from workstation software, and on some platforms the servers and workstation software are very similar. On the Macintosh operating system (OS) both the server and workstation are identical, not just functionally, but down to the code. The same installation disks will install the OS on the server and the workstation. It is the applications or server applications that make the difference.

This is effective because in a network solution, the Macintosh computer can be one of the most secure systems available. The server OS is identical to the workstation OS. The inherent design of the OS is for single users, not multiple users. There is no physical way to access information on a networked Macintosh computer, unless the computer has been specifically set up for this purpose. Even then, there are no backdoors, secret Apple passwords, or keystrokes to get around it. The operating system even has the ability to pass information through the system without touching the system in applications such as firewalls, e-mail servers, etc.

The Macintosh becomes available to multiple users when specific software like At Ease, AppleShare File Server and File Sharing is enabled.

Three basic software packages are used in schools today: FoolProof, MacJanet, and the Apple Network Administrator Toolkit. They all are intended to ease administration on the part of the teacher, but they all have various security features, advantages, and disadvantages. The Apple Network Administrator Toolkit includes the Network Assistant, At Ease Workstation, At Ease Server, and the User and Group Manager.

This appendix defines how to best implement the Macintosh in a school environment, using the Apple Network Administrator Toolkit. These software packages run on any network: LocalTalk, TokenTalk, EtherTalk, and TCP/IP. In addition, this software will run on any Macintosh with a hard drive. Supplementary information on FoolProof and MacJanet is included at the end of this appendix.

AT EASE

At Ease is currently the best workgroup and security software available for the Macintosh computer. Staff members with appropriate access can administer it from any location on the network.

At Ease provides three options for interfaces (all staff controlled). In the panel interface, instead of a hard drive, trash can, complex menus, etc., all the programs and documents the user has access to are large buttons on the screen. When the cursor is placed over each one, the computer verbalizes the title. Students do not need to know how to click-and-drag, open folders, learn hierarchical filing systems, etc. This works well for lower age groups and functionally challenged students. The reduced Finder gives the user the standard Macintosh interface, without the ability to access control panels, various system files, and other software. This interface is preferred for upper elementary students through staff members as it gives them all the features of using the Macintosh but prevents them from "exploring" and potentially corrupting the system. The full Finder, which should be available only to authorized staff, allows the user to customize control panels, network settings, and various other desktop features.

At Ease gives all three options, on any workstation, to the appropriate user. When the computer is activated, the user is prompted for a name and password. This information is passed to the At Ease Server which then sets up the environment for the user. A kindergarten student might get the simple button/panel interface when logging on while an older student or staff member gets a different interface on the same computer based on the privileges assigned their user ID. In addition, all of their own personal settings (i.e., Internet e-mail address, preferences, etc.) are set up on the workstation. A user can go to any computer in a networked school, and it will look the same, feel the same, and have the same access as any other computer.

Since this software is based on the At Ease Server, the user information is entered only once on the server, at which point the user can use any computer that is authorized by that server.

While not necessarily a file server, At Ease is designed to work in conjunction with AppleShare file servers to provide robust and secure file access. Students can not in any way access the files of staff members or other students on the server, while staff

members can access all students files, but not all staff files. Additionally, all passwords are encrypted when communicating over the network.

Appleshare File servers are not discussed in this paper. Their purpose is simply to store files and security is unquestioned.

SECURITY

One of the best-kept secrets of the Macintosh and of At Ease software is its security. The basic Macintosh OS is fully secure from network intrusion. There is no way to affect, modify, change, or view files on a networked Macintosh, unless the user expressly enables such capabilities. Only users physically sitting in front of the computer can access contents.

At Ease provides excellent desktop security. As previously mentioned, the user must log on to the network to access any resources on the computer. There is no way to access the hard drive if the user does not know the password. Someone trying to break in by using a startup disk is unable to access the hard drive because At Ease makes changes to the driver itself, rendering the hard drive unmountable and thus unviewable. All the user can do is reformat the hard drive which can be automatically rewritten.

At Ease maintains a list of permitted applications for each user. Although the hard drive may be filled with applications, some users will be able to launch only some programs. For example, lower elementary students could be allowed to only use ClarisWorks and other elementary programs. Older students may have permissions to use only ClarisWorks, Netscape, etc. This is a good way of locking out the Internet, in that students who violate the jurisdiction's acceptable use policy will be locked out of any browser software, but otherwise can use the computer. This also eliminates the ability to download various "hacking" software from the Internet. Students will be able to download these programs but unable to run them, as they will not be on the At Ease application list. Later in the day, when synchronization happens, all additional programs are erased.

Printing is another strong feature of the At Ease software. Individual users can have a printer or printers assigned to them so they can print only to their permitted printer. No matter which computer a student uses in the school, s/he still has access to the same printers. High speed office printers, colour laser printers, and any other printer can be made unavailable to any user. This also saves the user time in figuring out which printer they sent their job to.

Other At Ease security features prevent any user from making changes to the local system and applications. At Ease can be configured to allow saving to the local hard drive in either a shared folder or a user-specific folder. In the first case, any user can view the files, and in the second only the user can view/modify the files.

Lastly, if the workstation is disconnected from the network, the local At Ease software will still prevent access to the hard drive except by the administrator, or local users if so configured.

NETWORK ASSISTANT

Another tool in the Network Administrator Toolkit is the Network Assistant. This application saves an incredible amount of time and energy in configuring groups or labs of computers. It also allows teachers to teach using different tools, which keeps the students' attention and makes the teacher's job easier.

Teachers do not have the time to take a group of disks and physically install a piece of software on every single computer. This is one task that the Network Assistant simplifies. The staff member loads the software on a single computer, and the Network Assistant then synchronizes all the other computers so they are identical. Hit one button on the Network Assistant screen, and all computers are synchronized. This synchronization is done using broadcast packets, so fifty computers are synchronized in the same time it takes to set up one. These broadcast packets do not significantly slow down the network, making it unavailable to other users. Synchronization can be done over LocalTalk or EtherTalk networks, and TCP/IP. In addition, it can be timed to automatically take place after school hours.

Using the Network Assistant, staff can make instantaneous changes to all computers on the network. Settings like the date and time of each computer can be set instantly at any time. As well, staff can change the volume level of any or all computers while they are being used, change the screen resolution, and even make minor repairs like rebuilding the desktop. All computers can be set, within seconds, to use the same printer. However, the At Ease software, if implemented, will reset the printer for each individual user.

Support staff find the Network Assistant valuable. They can gain information about any workstation from anywhere on the wide area network. They can obtain reports on what applications are loaded on the hard drive, RAM, TCP/IP, and other configurations; available and used space on the hard drives; and many other functions. There is no need to check each computer's contents and configuration.

THE NETWORK ASSISTANT AS A TEACHING TOOL

Many features of the Network Assistant are also invaluable teaching aids for use in computer labs.

The Network Assistant allows teachers to share their screen with all the students or share any student's screen with every other student. Using this tool, the teacher can demonstrate skills and programs to all the students or show a student's work to everyone else without the expense of a projection system. In addition, the students can not override the screen sharing, so while the teacher is sharing a screen they can only observe.

The teacher also has the ability to lock any or all students' keyboards and screens. When their screen and keyboard are locked, students can not do any computer work and their attention can not be diverted to whatever is on the computer. The screen locking/unlocking is done instantly with the click of a button.

The teacher also can broadcast voice over the network to all the computers. This feature is useful when computers are in another room or when for some reason the teacher can not be heard. This voice broadcast can be done to a single computer or any number of computers simultaneously. Additionally, if the students' computers are equipped with microphones, the teacher can have a full two-way conversation with any student. Since teachers always have the ability to view or control any student workstation from their workstation, they can remotely help students with their work. Students also may request attention and/or help via their computers. This is useful in an examination situation where movement and other distracting noises would be unwanted.

Using the Network Assistant the teacher can remotely shut down or restart any computer or all the computers. At the end of the day, the teacher can shut down everything without going to every single workstation. The teacher also can watch what applications are being used at any time, remotely open or close any application that the students have access to, and instantly delete any file or application.

FOOLPROOF

FoolProof has long been used in Macintosh-based schools. In its original form, it was a simple tool for locking students out of the system folder and preventing tampering with various extensions and control panels. It also could "lock" the hard drive, preventing the user from saving any unwanted materials or making modifications to the hard drive. Recent versions of this software allow staff members to set FoolProof settings on one computer and have all the other computers automatically update their settings. This software was very effective before computers were connected to servers and the Internet.

In the original FoolProof, users could not make changes but applications could. For example, users could not create new folders when using the Finder, but if they were using an application (i.e., ClarisWorks) they could create as many folders as they wanted to. FoolProof also prevented dragging of files and folders in the Finder, but any application that had moving or deleting files as options within the application were not affected by FoolProof. While FoolProof locked the hard drive and prevented changes to the system folder, it could not prevent users from using any application on the hard drive. Although it is easy to ensure that only approved programs are on the hard drive initially, the advent of Internet web browsing software has allowed students to download new programs to the hard drive which could counteract the FoolProof software.

FoolProof continues to be one of the best ways of securing non-networked Macintosh computers in school systems today, where computers are connected to servers and the Internet.

MACJANET

MacJanet, one of the first "network resource administration" tools, makes unique features available to teachers, administrative staff, and students.

A MacJanet server controls all the printers, applications, and files on the server. The computer either has a hard drive with basic system software on it or no hard drive, but a bootable floppy disk. Running without a hard drive is very useful for labs where Mac Plus computers are used. All the lab needs is startup disks which allow Mac Plus computers to connect to the MacJanet server. Once connected, users can run applications and store files.

MacJanet allows the creation of many student and teacher volumes. Students see only their own files, which are accessible from any computer. Teachers see only their files and any student files. In addition, maximum storage spaces for each user can be set. For example, you could give each student 2 MB of storage space, and each teacher 10 MB. Neither would be allowed to use space beyond their limit.

MacJanet also allows for excellent printer control. Like At Ease, you can set limits on how many pages each user can print. However, unlike At Ease, you can not limit different users to different printers. MacJanet makes this possible by forcing all users to print through a MacJanet print queue. This has many advantages, including the ability to monitor which users print which jobs and how many pages they have sent to any printer. Also, people can continue working while printing occurs.

Originally, MacJanet allowed applications to be run from the server, thus negating the need for large hard drives on the workstation. However, with the advent of modern computer software, it is no longer feasible to run applications from the server. First, many applications require certain portions to be installed on the local hard drive. Second, network bandwidth prohibits users from launching simultaneous copies of applications over the network, as the network is bogged down almost immediately. Thirdly, many applications can not be run simultaneously by many different users.

MacJanet is not meant as security software over the local computer and hard drive. It is meant as a secure way of accessing and using file and printer services over the network.

MacJanet, while lacking many of the modern tools and features available to users of AppleShare Servers in conjunction with the Apple Network Administrator Toolkit, is still a viable option in a school of any size that requires robust and secure file services, good administrative capability over file servers, and good print volume monitoring.

SUMMARY

The Macintosh operating system, while intrinsically secure, can easily be used in a network environment, with tools like the Network Administrator Toolkit to create a truly user-friendly computer experience. Using the same operating system on both the workstation and the server makes it easy for staff members familiar with either one to administer and troubleshoot. With software like At Ease, FoolProof, or MacJanet, extensive desktop security is available, and performing functions like adding/deleting users is easy for staff members at all levels of expertise. The Network Administrator Toolkit transforms the Macintosh into a Network Operating System, with security, and administrative features comparable to any other platform.

APPENDIX E RELATED ALBERTA EDUCATION RESOURCES

Developing A Three-Year Technology Integration Plan: A Resource for School Jurisdictions (1999).

FOIPP and Technology: Best Practices for Alberta School Jurisdictions (1999).

FOIPP and Technology Highlights: Best Practices for Alberta School Jurisdictions (1999).

Implementing and Managing Web Site Development in Education: Best Practices for Alberta School Jurisdictions (1999).

Managing Technology Funding: Best Practices for Alberta School Jurisdictions (1999).

Network Design: Best Practices for Alberta School Jurisdictions (1999).

On-Line Learning: Best Practices for Alberta School Jurisdictions (1999).

Preparing to Implement Learner Outcomes in Technology: Best Practices for Alberta School Jurisdictions (1999).

Professional Development for Teaching Technology Across the Curriculum: Best Practices for Alberta School Jurisdictions (1999).

Technical Support Planning: Best Practices for Alberta School Jurisdictions (1999).

Technology Implementation Review, Grande Yellowhead Regional Division No. 24 and Wolf Creek Regional Division No. 32: Best Practices and Key Learnings with Respect to Technology, Its Implementation and Management in Education (1997).



U.S. Department of Education
Office of Educational Research and Improvement (OERI)
National Library of Education (NLE)
Educational Resources Information Center (ERIC)



NOTICE

REPRODUCTION BASIS



This document is covered by a signed "Reproduction Release (Blanket) form (on file within the ERIC system), encompassing all or classes of documents from its source organization and, therefore, does not require a "Specific Document" Release form.



This document is Federally-funded, or carries its own permission to reproduce, or is otherwise in the public domain and, therefore, may be reproduced by ERIC without a signed Reproduction Release form (either "Specific Document" or "Blanket").